ARM® Compiler toolchain v4.1 for µVision

Compiler Reference

Copyright © 2007-2008, 2011 ARM. All rights reserved.
ARM DUI 0376C (ID061811)

ARM Compiler toolchain v4.1 for µVision
Compiler Reference

Copyright © 2007-2008, 2011 ARM. All rights reserved.

Release Information

The following changes have been made to this book.

Change History

Date	Issue	Confidentiality	Change
May 2007	A	Non-Confidential	Release for RVCT v3.1 for µVision
December 2008	B	Non-Confidential	Release for RVCT v4.0 for µVision
June 2011	C	Non-Confidential	Release for ARM Compiler toolchain v4.1 for µVision

Proprietary Notice

Words and logos marked with or are registered trademarks or trademarks of ARM in the EU and other countries, except as otherwise stated below in this proprietary notice. Other brands and names mentioned herein may be the trademarks of their respective owners.

Neither the whole nor any part of the information contained in, or the product described in, this document may be adapted or reproduced in any material form except with the prior written permission of the copyright holder.

The product described in this document is subject to continuous developments and improvements. All particulars of the product and its use contained in this document are given by ARM in good faith. However, all warranties implied or expressed, including but not limited to implied warranties of merchantability, or fitness for purpose, are excluded.

This document is intended only to assist the reader in the use of the product. ARM shall not be liable for any loss or damage arising from the use of any information in this document, or any error or omission in such information, or any incorrect use of the product.

Where the term ARM is used it means “ARM or any of its subsidiaries as appropriate”.

Some material in this document is based on IEEE 754 - 1985 IEEE Standard for Binary Floating-Point Arithmetic. The IEEE disclaims any responsibility or liability resulting from the placement and use in the described manner.

Confidentiality Status

This document is Non-Confidential. The right to use, copy and disclose this document may be subject to license restrictions in accordance with the terms of the agreement entered into by ARM and the party that ARM delivered this document to.

Product Status

The information in this document is final, that is for a developed product.

Web Address

`http://www.arm.com`

Contents
ARM Compiler toolchain v4.1 for µVision Compiler Reference

Chapter 1 **Conventions and Feedback**

Chapter 2 **Introduction**

2.1 About the ARM compiler ... 2-2
2.2 Source language modes ... 2-3
2.3 Language extensions and language compliance ... 2-5
2.4 The C and C++ libraries ... 2-7

Chapter 3 **Compiler Command-line Options**

3.1 Command-line options ... 3-6

Chapter 4 **Language Extensions**

4.1 Preprocessor extensions ... 4-2
4.2 C99 language features available in C90 ... 4-4
4.3 C99 language features available in C++ and C90 ... 4-6
4.4 Standard C language extensions ... 4-9
4.5 Standard C++ language extensions ... 4-13
4.6 Standard C and Standard C++ language extensions ... 4-16

Chapter 5 **Compiler-specific Features**

5.1 Keywords and operators ... 5-2
5.2 __declspec attributes ... 5-19
5.3 Function attributes ... 5-25
5.4 Type attributes ... 5-36
5.5 Variable attributes ... 5-39
5.6 Pragmas ... 5-47

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.7 Instruction intrinsics 5-61
5.8 ARMv6 SIMD intrinsics 5-88
5.9 ETSI basic operations 5-89
5.10 C55x intrinsics 5-91
5.11 VFP status intrinsic 5-92
5.12 Fused Multiply Add (FMA) intrinsics 5-93
5.13 Named register variables 5-94
5.14 Compiler predefines 5-98

Chapter 6 C and C++ Implementation Details

6.1 C and C++ implementation details 6-2
6.2 C++ implementation details 6-11

Appendix A ARMv6 SIMD Instruction Intrinsics

A.1 ARMv6 SIMD intrinsics by prefix A-3
A.2 ARMv6 SIMD intrinsics, summary descriptions, byte lanes, affected flags A-5
A.3 ARMv6 SIMD intrinsics, compatible processors and architectures A-9
A.4 ARMv6 SIMD instruction intrinsics and APSR GE flags A-10
A.5 __qadd16 intrinsic A-11
A.6 __qadd8 intrinsic A-12
A.7 __qasx intrinsic A-13
A.8 __qsax intrinsic A-14
A.9 __qsub16 intrinsic A-15
A.10 __qsub8 intrinsic A-16
A.11 __sadd16 intrinsic A-17
A.12 __sadd8 intrinsic A-18
A.13 __sasx intrinsic A-19
A.14 __sel intrinsic A-20
A.15 __shadd16 intrinsic A-21
A.16 __shadd8 intrinsic A-22
A.17 __shasx intrinsic A-23
A.18 __shsax intrinsic A-24
A.19 __shsub16 intrinsic A-25
A.20 __shsub8 intrinsic A-26
A.21 __smlad intrinsic A-27
A.22 __smladx intrinsic A-28
A.23 __smlald intrinsic A-29
A.24 __smlaldx intrinsic A-30
A.25 __smlsd intrinsic A-31
A.26 __smlsdx intrinsic A-32
A.27 __smlsld intrinsic A-33
A.28 __smlsldx intrinsic A-34
A.29 __smuad intrinsic A-35
A.30 __smusd intrinsic A-36
A.31 __smusdx intrinsic A-37
A.32 __smuadx intrinsic A-38
A.33 __ssat16 intrinsic A-39
A.34 __ssax intrinsic A-40
A.35 __ssub16 intrinsic A-41
A.36 __ssub8 intrinsic A-42
A.37 __sxtab16 intrinsic A-43
A.38 __sxtb16 intrinsic A-44
A.39 __uadd16 intrinsic A-45
A.40 __uadd8 intrinsic A-46
A.41 __uasx intrinsic A-47
A.42 __uhadd16 intrinsic A-48
A.43 __uhadd8 intrinsic A-49
A.44 __uhasx intrinsic A-50
A.45 __uhsax intrinsic A-51
A.46 __uhsub16 intrinsic A-52

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.47 __uhsub8 intrinsic A-53
A.48 __uqadd16 intrinsic A-54
A.49 __uqadd8 intrinsic A-55
A.50 __uqasx intrinsic A-56
A.51 __uqsax intrinsic A-57
A.52 __uqsub16 intrinsic A-58
A.53 __uqsub8 intrinsic A-59
A.54 __usad8 intrinsic A-60
A.55 __usada8 intrinsic A-61
A.56 __usax intrinsic A-62
A.57 __usat16 intrinsic A-63
A.58 __usub16 intrinsic A-64
A.59 __usub8 intrinsic A-65
A.60 __uxtab16 intrinsic A-66
A.61 __uxtb16 intrinsic A-67

Appendix B Via File Syntax
B.1 Overview of via files B-2
B.2 Syntax B-3

Appendix C Summary Table of GNU Language Extensions

Appendix D Standard C Implementation Definition
D.1 Implementation definition D-2
D.2 Behaviors considered undefined by the ISO C Standard D-8

Appendix E Standard C++ Implementation Definition
E.1 Integral conversion E-2
E.2 Calling a pure virtual function E-3
E.3 Major features of language support E-4
E.4 Standard C++ library implementation definition E-5

Appendix F C and C++ Compiler Implementation Limits
F.1 C++ ISO/IEC standard limits F-2
F.2 Limits for integral numbers F-4
F.3 Limits for floating-point numbers F-5

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Chapter 1
Conventions and Feedback

The following describes the typographical conventions and how to give feedback:

Typographical conventions

The following typographical conventions are used:

`monospace` Denotes text that can be entered at the keyboard, such as commands, file and program names, and source code.

`monospace` Denotes a permitted abbreviation for a command or option. The underlined text can be entered instead of the full command or option name.

`monospace italic`

Denotes arguments to commands and functions where the argument is to be replaced by a specific value.

`monospace bold`

Denotes language keywords when used outside example code.

italic Highlights important notes, introduces special terminology, denotes internal cross-references, and citations.

bold Highlights interface elements, such as menu names. Also used for emphasis in descriptive lists, where appropriate, and for ARM® processor signal names.

Feedback on this product

If you have any comments and suggestions about this product, contact your supplier and give:

- your name and company

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- the serial number of the product
- details of the release you are using
- details of the platform you are using, such as the hardware platform, operating system type and version
- a small standalone sample of code that reproduces the problem
- a clear explanation of what you expected to happen, and what actually happened
- the commands you used, including any command-line options
- sample output illustrating the problem
- the version string of the tools, including the version number and build numbers.

Feedback on documentation

If you have comments on the documentation, e-mail `errata@arm.com`. Give:

- the title
- the number, ARM DUI 0376C
- if viewing online, the topic names to which your comments apply
- if viewing a PDF version of a document, the page numbers to which your comments apply
- a concise explanation of your comments.

ARM also welcomes general suggestions for additions and improvements.

ARM periodically provides updates and corrections to its documentation on the ARM Information Center, together with knowledge articles and *Frequently Asked Questions* (FAQs).

Other information

- ARM Product Manuals, `http://www.keil.com/support/man_arm.htm`
- Keil Support Knowledgebase, `http://www.keil.com/support/knowledgebase.asp`
- Keil Product Support, `http://www.keil.com/support/`
- ARM Glossary, `http://infocenter.arm.com/help/topic/com.arm.doc.aeg0014-/index.html`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Chapter 2
Introduction

The following topics introduce the compiler, `armcc`:

- *About the ARM compiler* on page 2-2
- *Source language modes* on page 2-3
- *Language extensions and language compliance* on page 2-5
- *The C and C++ libraries* on page 2-7.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

2.1 About the ARM compiler

The compiler, `armcc`, enables you to compile your C and C++ code.

The compiler:

- Is an optimizing compiler. Command-line options enable you to control the level of optimization.
- Compiles:
  - ISO Standard C:1990 source
  - ISO Standard C:1999 source
  - ISO Standard C++:2003 source

  into:
  - 32-bit ARM code
  - 16/32-bit Thumb-2 code
  - 16-bit Thumb code.
- Complies with the *Base Standard Application Binary Interface for the ARM Architecture (*BSABI). In particular, the compiler:
  - Generates output objects in ELF format.
  - Generates DWARF Debugging Standard Version 3 (DWARF 3) debug information and contains support for DWARF 2 debug tables.

  See *Compliance with the Application Binary Interface (ABI) for the ARM architecture* on page 2-9 in *Using ARM® C and C++ Libraries and Floating-Point Support* for more information.
- Can generate an assembly language listing of the output code, and can interleave an assembly language listing with source code.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

2.2 Source language modes

The compiler has three distinct source language modes that you can use to compile different varieties of C and C++ source code. These are:

- ISO C90
- ISO C99
- ISO C++.

2.2.1 ISO C90

The compiler compiles C as defined by the 1990 C standard and addenda:

- ISO/IEC 9899:1990. The 1990 International Standard for C.
- ISO/IEC 9899 AM1. The 1995 Normative Addendum 1, adding international character support through `wchar.h` and `wtype.h`.

The compiler also supports several extensions to ISO C90. See *Language extensions and language compliance* on page 2-5 for more information.

Throughout this document, the term:

C90 Means ISO C90, together with the ARM extensions.

Use the compiler option `--c90` to compile C90 code. This is the default.

Strict C90 Means C as defined by the 1990 C standard and addenda.

See also

- *--c90* on page 3-18
- *--strict, --no_strict* on page 3-88
- *Language extensions and language compliance* on page 2-5
- Appendix D *Standard C Implementation Definition*.

2.2.2 ISO C99

The compiler compiles C as defined by the 1999 C standard and addenda:

- ISO/IEC 9899:1999. The 1999 International Standard for C.

The compiler also supports several extensions to ISO C99. See *Language extensions and language compliance* on page 2-5 for more information.

Throughout this document, the term:

C99 Means ISO C99, together with the ARM extensions.

Use the compiler option `--c99` to compile C99 code.

Strict C99 Means C as defined by the 1999 C standard and addenda.

Standard C Means C90 or C99 as appropriate.

C Means any of C90, strict C90, C99, and Standard C.

See also

- *--c99* on page 3-18
- *--strict, --no_strict* on page 3-88
- *Language extensions and language compliance* on page 2-5

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- Appendix D *Standard C Implementation Definition*.

2.2.3 ISO C++

The compiler compiles C++ as defined by the 2003 standard, excepting wide streams and export templates:

- ISO/IEC 14822:2003. The 2003 International Standard for C++.

The compiler also supports several extensions to ISO C++. See *Language extensions and language compliance* on page 2-5 for more information.

Throughout this document, the term:

strict C++ Means ISO C++, excepting wide streams and export templates.

Standard C++ Means strict C++.

C++ Means ISO C++, excepting wide streams and export templates, either with or without the ARM extensions.

Use the compiler option `--cpp` to compile C++ code.

See also

- *--cpp* on page 3-20
- *--strict, --no_strict* on page 3-88
- *Language extensions and language compliance* on page 2-5
- Appendix E *Standard C++ Implementation Definition*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

2.3 Language extensions and language compliance

The compiler supports numerous extensions to its various source languages. It also provides several command-line options for controlling compliance with the available source languages.

2.3.1 Language extensions

The language extensions supported by the compiler are categorized as follows:

C99 features

The compiler makes some language features of C99 available:

- as extensions to strict C90, for example, `//`-style comments
- as extensions to both Standard C++ and strict C90, for example, **`restrict`** pointers.

For more information see:

- *C99 language features available in C90* on page 4-4
- *C99 language features available in C++ and C90* on page 4-6.

Standard C extensions

The compiler supports numerous extensions to strict C99, for example, function prototypes that override old-style nonprototype definitions. See *Standard C language extensions* on page 4-9 for more information.

These extensions to Standard C are also available in C90.

Standard C++ extensions

The compiler supports numerous extensions to strict C++, for example, qualified names in the declaration of class members. See *Standard C++ language extensions* on page 4-13 for more information.

These extensions are not available in either Standard C or C90.

Standard C and Standard C++ extensions

The compiler supports some extensions specific to strict C++ and strict C90, for example, anonymous classes, structures, and unions. See *Standard C and Standard C++ language extensions* on page 4-16 for more information.

ARM-specific extensions

The compiler supports a range of extensions specific to the ARM compiler, for example, instruction intrinsics and other builtin functions. See Chapter 5 *Compiler-specific Features* for more information.

2.3.2 Language compliance

The compiler has a mode where compliance to a source language is enforced.

In strict mode the compiler enforces compliance with the language standard relevant to the source language. For example, the use of //-style comments results in an error when compiling strict C90.

To compile in strict mode, use the command-line option `--strict`.

For example, compiling a `.cpp` file with the command-line option `--strict` compiles Standard C++

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--strict, --no_strict* on page 3-88
- *Filename suffixes recognized by the compiler* on page 3-14 in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

2.4 The C and C++ libraries

The following runtime C and C++ libraries are provided:

The ARM C libraries

The ARM C libraries provide standard C functions, and helper functions used by the C and C++ libraries. The C libraries also provide target-dependent functions that are used to implement the standard C library functions such as `printf` in a semihosted environment. The C libraries are structured so that you can redefine target-dependent functions in your own code to remove semihosting dependencies.

The ARM libraries comply with:

- the *C Library ABI for the ARM Architecture* (CLIBABI)
- the *C++ ABI for the ARM Architecture* (CPPABI).

See *Compliance with the Application Binary Interface (ABI) for the ARM architecture* on page 2-9 in *Using ARM® C and C++ Libraries and Floating-Point Support* for more information.

Rogue Wave Standard C++ Library version 2.02.03

The Rogue Wave Standard C++ Library, as supplied by Rogue Wave Software, Inc., provides Standard C++ functions and objects such as `cout`. It includes data structures and algorithms known as the *Standard Template Library* (STL). The C++ libraries use the C libraries to provide target-specific support. The Rogue Wave Standard C++ Library is provided with C++ exceptions enabled.

For more information on the Rogue Wave libraries, see the Rogue Wave web site at: `http://www.roguewave.com`

Support libraries

The ARM C libraries provide additional components to enable support for C++ and to compile code for different architectures and processors.

The C and C++ libraries are provided as binaries only. There is a variant of the 1990 ISO Standard C library for each combination of major build options, such as the byte order of the target system, whether interworking is selected, and whether floating-point support is selected.

See Chapter 2 *The ARM C and C++ libraries* in *Using ARM® C and C++ Libraries and Floating-Point Support* for more information.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Chapter 3
Compiler Command-line Options

This chapter lists the command-line options accepted by the compiler, `armcc`. The options are:

- *-Aopt* on page 3-6
- *--allow_null_this, --no_allow_null_this* on page 3-6
- *--alternative_tokens, --no_alternative_tokens* on page 3-7
- *--anachronisms, --no_anachronisms* on page 3-7
- *--apcs=qualifer...qualifier* on page 3-7
- *--arm* on page 3-11
- *--arm_only* on page 3-11
- *--asm* on page 3-12
- *--asm_dir=directory_name* on page 3-13
- *--autoinline, --no_autoinline* on page 3-14
- *--bigend* on page 3-14
- *--bitband* on page 3-15
- *--brief_diagnostics, --no_brief_diagnostics* on page 3-15
- *--bss_threshold=num* on page 3-16
- *-c* on page 3-17
- *-C* on page 3-17
- *--c90* on page 3-18
- *--c99* on page 3-18
- *--code_gen, --no_code_gen* on page 3-18
- *--compatible=name* on page 3-19
- *--compile_all_input, --no_compile_all_input* on page 3-20
- *--cpp* on page 3-20

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *--cpu=list* on page 3-20
- *--cpu=name* on page 3-20
- *--create_pch=filename* on page 3-22
- *-Dname[(parm-list)][=def]* on page 3-22
- *--data_reorder, --no_data_reorder* on page 3-23
- *--debug, --no_debug* on page 3-24
- *--debug_macros, --no_debug_macros* on page 3-24
- *--default_definition_visibility=visibility* on page 3-24
- *--default_extension=ext* on page 3-25
- *--dep_name, --no_dep_name* on page 3-25
- *--depend=filename* on page 3-26
- *--depend_dir=directory_name* on page 3-27
- *--depend_format=string* on page 3-27
- *--depend_single_line, --no_depend_single_line* on page 3-28
- *--depend_system_headers, --no_depend_system_headers* on page 3-29
- *--depend_target=target* on page 3-30
- *--diag_error=tag[,tag,...]* on page 3-30
- *--diag_remark=tag[,tag,...]* on page 3-31
- *--diag_style={arm|ide|gnu}* on page 3-31
- *--diag_suppress=tag[,tag,...]* on page 3-32
- *--diag_suppress=optimizations* on page 3-33
- *--diag_warning=tag[,tag,...]* on page 3-33
- *--diag_warning=optimizations* on page 3-34
- *--dollar, --no_dollar* on page 3-34
- *--dwarf2* on page 3-35
- *--dwarf3* on page 3-35
- *-E* on page 3-35
- *--emit_frame_directives, --no_emit_frame_directives* on page 3-36
- *--enum_is_int* on page 3-36
- *--errors=filename* on page 3-36
- *--exceptions, --no_exceptions* on page 3-37
- *--exceptions_unwind, --no_exceptions_unwind* on page 3-37
- *--export_all_vtbl, --no_export_all_vtbl* on page 3-38
- *--export_defs_implicitly, --no_export_defs_implicitly* on page 3-38
- *--extended_initializers, --no_extended_initializers* on page 3-38
- *--feedback=filename* on page 3-39
- *--force_new_nothrow, --no_force_new_nothrow* on page 3-40
- *--forceinline* on page 3-40
- *--fp16_format=format* on page 3-41
- *--fpmode=model* on page 3-42
- *--fpu=list* on page 3-43
- *--fpu=name* on page 3-44
- *--friend_injection, --no_friend_injection* on page 3-46
- *-g* on page 3-47
- *--global_reg=reg_name[,reg_name,...]* on page 3-47
- *--guiding_decls, --no_guiding_decls* on page 3-48
- *--help* on page 3-49

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *--hide_all, --no_hide_all* on page 3-49
- *-Idir[,dir,...]* on page 3-49
- *--ignore_missing_headers* on page 3-50
- *--implicit_include, --no_implicit_include* on page 3-50
- *--implicit_include_searches, --no_implicit_include_searches* on page 3-51
- *--implicit_key_function, --no_implicit_key_function* on page 3-51
- *--implicit_typename, --no_implicit_typename* on page 3-52
- *--import_all_vtbl* on page 3-52
- *--info=totals* on page 3-52
- *--inline, --no_inline* on page 3-53
- *--interface_enums_are_32_bit* on page 3-53
- *--interleave* on page 3-54
- *-Jdir[,dir,...]* on page 3-55
- *--kandr_include* on page 3-55
- *-Lopt* on page 3-56
- *--library_interface=lib* on page 3-56
- *--library_type=lib* on page 3-58
- *--link_all_input, --no_link_all_input* on page 3-58
- *--list* on page 3-59
- *--list_dir=directory_name* on page 3-61
- *--list_macros* on page 3-61
- *--littleend* on page 3-61
- *--locale=lang_country* on page 3-62
- *--long_long* on page 3-62
- *--loose_implicit_cast* on page 3-63
- *--lower_ropi, --no_lower_ropi* on page 3-63
- *--lower_rwpi, --no_lower_rwpi* on page 3-63
- *--ltcg* on page 3-64
- *-M* on page 3-64
- *--md* on page 3-65
- *--message_locale=lang_country[.codepage]* on page 3-65
- *--min_array_alignment=opt* on page 3-66
- *--mm* on page 3-67
- *--multibyte_chars, --no_multibyte_chars* on page 3-67
- *--multifile, --no_multifile* on page 3-67
- *--multiply_latency=cycles* on page 3-68
- *--narrow_volatile_bitfields* on page 3-69
- *--nonstd_qualifier_deduction, --no_nonstd_qualifier_deduction* on page 3-69
- *-o filename* on page 3-69
- *-Onum* on page 3-71
- *--old_specializations, --no_old_specializations* on page 3-72
- *--old_style_preprocessing* on page 3-72
- *-Ospace* on page 3-72
- *-Otime* on page 3-73
- *--output_dir=directory_name* on page 3-73
- *-P* on page 3-74
- *--parse_templates, --no_parse_templates* on page 3-74

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *--pch* on page 3-75
- *--pch_dir=dir* on page 3-75
- *--pch_messages, --no_pch_messages* on page 3-76
- *--pch_verbose, --no_pch_verbose* on page 3-76
- *--pending_instantiations=n* on page 3-76
- *--phony_targets* on page 3-77
- *--pointer_alignment=num* on page 3-77
- *--preinclude=filename* on page 3-78
- *--preprocess_assembly* on page 3-79
- *--preprocessed* on page 3-80
- *--reassociate_saturation, --no_reassociate_saturation* on page 3-80
- *--reduce_paths, --no_reduce_paths* on page 3-81
- *--relaxed_ref_def, --no_relaxed_ref_def* on page 3-82
- *--remarks* on page 3-82
- *--remove_unneeded_entities, --no_remove_unneeded_entities* on page 3-82
- *--restrict, --no_restrict* on page 3-83
- *--retain=option* on page 3-83
- *--rtti, --no_rtti* on page 3-84
- *--rtti_data, --no_rtti_data* on page 3-84
- *-S* on page 3-85
- *--show_cmdline* on page 3-85
- *--signed_bitfields, --unsigned_bitfields* on page 3-86
- *--signed_chars, --unsigned_chars* on page 3-86
- *--split_ldm* on page 3-87
- *--split_sections* on page 3-87
- *--strict, --no_strict* on page 3-88
- *--strict_warnings* on page 3-89
- *--sys_include* on page 3-90
- *--thumb* on page 3-90
- *--trigraphs, --no_trigraphs* on page 3-91
- *--type_traits_helpers, --no_type_traits_helpers* on page 3-91
- *-Uname* on page 3-91
- *--unaligned_access, --no_unaligned_access* on page 3-92
- *--use_frame_pointer* on page 3-93
- *--use_pch=filename* on page 3-93
- *--using_std, --no_using_std* on page 3-94
- *--version_number* on page 3-94
- *--vfe, --no_vfe* on page 3-94
- *--via=filename* on page 3-95
- *--visibility_inlines_hidden* on page 3-95
- *--vla, --no_vla* on page 3-96
- *--vsn* on page 3-96
- *-W* on page 3-96
- *--wchar, --no_wchar* on page 3-97
- *--wchar16* on page 3-97
- *--wchar32* on page 3-97
- *--whole_program* on page 3-98

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *--wrap_diagnostics, --no_wrap_diagnostics* on page 3-98.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1 Command-line options

This section lists the command-line options supported by the compiler in alphabetical order.

3.1.1 `-Aopt`

This option specifies command-line options to pass to the assembler when it is invoked by the compiler to assemble either `.s` input files or embedded assembly language functions.

Syntax

`-Aopt`

Where:

`opt` is a command-line option to pass to the assembler.

Note

Some compiler command-line options are passed to the assembler automatically whenever it is invoked by the compiler. For example, if the option `--cpu` is specified on the compiler command line, then this option is passed to the assembler whenever it is invoked to assemble `.s` files or embedded assembler.

To see the compiler command-line options passed by the compiler to the assembler, use the compiler command-line option `-A--show_cmdline`.

Example

```
armcc -A--predefine="NEWVERSION SETL {TRUE}" main.c
```

Restrictions

If an unsupported option is passed through using `-A`, an error is generated by the assembler.

See also

- *--cpu=name* on page 3-20
- *-Lopt* on page 3-56
- *--show_cmdline* on page 3-85.

3.1.2 `--allow_null_this, --no_allow_null_this`

These options allow and disallow null **`this`** pointers in C++.

Usage

Allowing null **`this`** pointers gives well-defined behavior when a nonvirtual member function is called on a null object pointer.

Disallowing null **`this`** pointers enables the compiler to perform optimizations, and conforms with the C++ standard.

Default

The default is `--no_allow_null_this`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.3 `--alternative_tokens`, `--no_alternative_tokens`

This option enables or disables the recognition of alternative tokens in C and C++.

Usage

In C and C++, use this option to control recognition of the digraphs. In C++, use this option to control recognition of operator keywords, for example, **`and`** and **`bitand`**.

Default

The default is `--alternative_tokens`.

3.1.4 `--anachronisms`, `--no_anachronisms`

This option enables or disables anachronisms in C++.

Mode

This option is effective only if the source language is C++.

Default

The default is `--no_anachronisms`.

Example

```
typedef enum { red, white, blue } tricolor;
inline tricolor operator++(tricolor c, int)
{
    int i = static_cast<int>(c) + 1;
    return static_cast<tricolor>(i);
}
void foo(void)
{
    tricolor c = red;
    c++; // okay
    ++c; // anachronism
}
```

Compiling this code with the option `--anachronisms` generates a warning message.

Compiling this code without the option `--anachronisms` generates an error message.

See also

- *--cpp* on page 3-20
- *--strict, --no_strict* on page 3-88
- *--strict_warnings* on page 3-89
- *Anachronisms* on page 6-11.

3.1.5 `--apcs=`*`qualifer...qualifier`*

This option controls interworking and position independence when generating code.

By specifying qualifiers to the `--apcs` command-line option, you can define the variant of the *Procedure Call Standard for the ARM architecture* (AAPCS) used by the compiler.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Syntax

`--apcs=qualifer...qualifier`

Where `qualifier...qualifier` denotes a list of qualifiers. There must be:

- at least one qualifier present
- no spaces separating individual qualifiers in the list.

Each instance of `qualifier` must be one of:

/interwork, /nointerwork
: Generates code with or without ARM/Thumb™ interworking support. The default is `/nointerwork`, except for ARMv5T and later where the default is `/interwork`.

/ropi, /noropi
: Enables or disables the generation of *Read-Only Position-Independent* (ROPI) code. The default is `/noropi`.

  `/[no]pic` is an alias for `/[no]ropi`.

/rwpi, /norwpi
: Enables or disables the generation of *Read/Write Position-Independent* (RWPI) code. The default is `/norwpi`.

  `/[no]pid` is an alias for `/[no]rwpi`.

/fpic, /nofpic
: Enables or disables the generation of read-only position-independent code where relative address references are independent of the location where your program is loaded.

/hardfp, /softfp
: Requests hardware or software floating-point linkage. This enables the procedure call standard to be specified separately from the version of the floating-point hardware available through the `--fpu` option. It is still possible to specify the procedure call standard by using the `--fpu` option, but the use of `--apcs` is recommended.

Note

You can alternatively specify multiple qualifiers. For example, `--apcs=/nointerwork/noropi/norwpi` is equivalent to `--apcs=/nointerwork --apcs=noropi/norwpi`.

Default

If you do not specify an `--apcs` option, the compiler assumes `--apcs=/nointerwork/noropi/norwpi/nofpic`.

Usage

`/interwork`, `/nointerwork`

By default, code is generated:

- without interworking support, that is `/nointerwork`, unless you specify a `--cpu` option that corresponds to architecture ARMv5T or later
- with interworking support, that is `/interwork`, on ARMv5T and later. ARMv5T and later architectures provide direct support to interworking by using instructions such as `BLX` and load to program counter instructions.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

`/ropi`, `/noropi`
: If you select the `/ropi` qualifier to generate ROPI code, the compiler:

- addresses read-only code and data PC-relative
- sets the *Position Independent* (PI) attribute on read-only output sections.

——— **Note** ———

`--apcs=/ropi` is not supported when compiling C++.

———————————

`/rwpi`, `/norwpi`
: If you select the `/rwpi` qualifier to generate RWPI code, the compiler:

- addresses writable data using offsets from the static base register `sb`. This means that:
  - the base address of the RW data region can be fixed at runtime
  - data can have multiple instances
  - data can be, but does not have to be, position-independent.
- sets the PI attribute on read/write output sections.

——— **Note** ———

Because the `--lower_rwpi` option is the default, code that is not RWPI is automatically transformed into equivalent code that is RWPI. This static initialization is done at runtime by the C++ constructor mechanism, even for C.

———————————

`/fpic`, `/nofpic`
: If you select this option, the compiler:

- accesses all static data using PC-relative addressing
- accesses all imported or exported read-write data using a *Global Offset Table* (GOT) entry created by the linker
- accesses all read-only data relative to the PC.

You must compile your code with `/fpic` if it uses shared objects. This is because relative addressing is only implemented when your code makes use of System V shared libraries.

You do not have to compile with `/fpic` if you are building either a static image or static library.

The use of `/fpic` is supported when compiling C++. In this case, virtual function tables and `typeinfo` are placed in read-write areas so that they can be accessed relative to the location of the PC.

——— **Note** ———

When building a System V shared library, use `--apcs /fpic` together with `--no_hide_all`.

———————————

`/hardfp`
: If you use `/hardfp`, the compiler generates code for hardware floating-point linkage. Hardware floating-point linkage uses the FPU registers to pass the arguments and return values.

`/hardfp` interacts with or overrides explicit or implicit use of `--fpu` as follows:

- If floating-point support is not permitted (for example, because `--fpu=none` is specified, or because of other means), `/hardfp` is ignored.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- If floating-point support is permitted, but without floating-point hardware (`--fpu=softvfp`), `/hardfp` gives an error.
- If floating-point hardware is available and the *hardfp* calling convention is used (`--fpu=vfp...`), `/hardfp` is ignored.
- If floating-point hardware is present and the *softfp* calling convention is used (`--fpu=softvfp+vfp...`), `/hardfp` gives an error.

The `/hardfp` and `/softfp` qualifiers are mutually exclusive.

`/softfp` If you use `/softfp`, software floating-point linkage is used. Software floating-point linkage means that the parameters and return value for a function are passed using the ARM integer registers `r0` to `r3` and the stack.

`/softfp` interacts with or overrides explicit or implicit use of `--fpu` as follows:

- If floating-point support is not permitted (for example, because `--fpu=none` is specified, or because of other means), `/softfp` is ignored.
- If floating-point support is permitted, but without floating-point hardware (`--fpu=softvfp`), `/softfp` is ignored because the state is already `/softfp`.
- If floating-point hardware is present, `/softfp` forces the *softfp* (`--fpu=softvfp+vfp...`) calling convention.

The `/hardfp` and `/softfp` qualifiers are mutually exclusive.

Restrictions

There are restrictions when you compile code with `/ropi`, or `/rwpi`, or `/fpic`.

`/ropi` The main restrictions when compiling with `/ropi` are:

- The use of `--apcs=/ropi` is not supported when compiling C++. You can compile only the C subset of C++ with `/ropi`.
- Some constructs that are legal C do not work when compiled for `--apcs=/ropi`. For example:

```
extern const int ci; // ro
const int *p2 = &ci; // this static initialization
                     // does not work with --apcs=/ropi
```

To enable such static initializations to work, compile your code using the `--lower_ropi` option. For example:

```
armcc --apcs=/ropi --lower_ropi
```

`/rwpi` The main restrictions when compiling with `/rwpi` are:

- Some constructs that are legal C do not work when compiled for `--apcs=/rwpi`. For example:

```
int i;               // rw
int *p1 = &i;        // this static initialization
                     // does not work with --apcs=/rwpi
                     // --no_lower_rwpi
```

To enable such static initializations to work, compile your code using the `--lower_rwpi` option. For example:

```
armcc --apcs=/rwpi
```

Note

You do not have to specify `--lower_rwpi`, because this is the default.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

`/fpic` The main restrictions when compiling with `/fpic` are:

- By default, if you use `--apcs=/fpic`, the compiler exports only functions and data marked `__declspec(dllexport)`.
- If you use `--apcs=/fpic` and `--no_hide_all` on the same command line, the compiler uses default ELF dynamic visibility for all `extern` variables and functions that do not use `__declspec(dll*)`. The compiler disables auto-inlining for functions with default ELF visibility.

See also

- *--fpu=name* on page 3-44
- *--hide_all, --no_hide_all* on page 3-49
- *--lower_ropi, --no_lower_ropi* on page 3-63
- *--lower_rwpi, --no_lower_rwpi* on page 3-63
- *__declspec(dllexport)* on page 5-19
- *Compiler options for floating-point linkage and computations on page 5-65*
- *ARM C libraries and multithreading* on page 2-16 in *Using ARM® C and C++ Libraries and Floating-Point Support*
- *Overview of veneers* on page 4-26 in *Using the ARM Linker*

3.1.6 `--arm`

This option is a request to the compiler to target the ARM instruction set. The compiler is permitted to generate both ARM and Thumb code, but recognizes that ARM code is preferred.

Note

This option is not relevant for Thumb-only processors such as Cortex-M4, Cortex-M3, Cortex-M1, and Cortex-M0.

Default

This is the default option for targets supporting the ARM instruction set.

See also

- *--arm_only*
- *--cpu=list* on page 3-20
- *--cpu=name* on page 3-20
- *--thumb* on page 3-90
- *#pragma arm* on page 5-47
- *ARM architectures supported by the toolchain* on page 2-14 in *Getting Started*.

3.1.7 `--arm_only`

This option enforces ARM-only code. The compiler behaves as if Thumb is absent from the target architecture.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

The compiler propagates the `--arm_only` option to the assembler and the linker.

Default

For targets that support the ARM instruction set, the default is `--arm`. For targets that do not support the ARM instruction set, the default is `--thumb`.

Example

```
armcc --arm_only myprog.c
```

Note

If you specify `armcc --arm_only --thumb myprog.c`, this does *not* mean that the compiler checks your code to ensure that no Thumb code is present. It means that `--thumb` overrides `--arm_only`, because of command-line ordering.

See also

- *--arm* on page 3-11
- *--thumb* on page 3-90
- *Assembler command line options* on page 2-3 in the *Assembler Reference* for information on `--16` and `--32`
- *About ordering the compilation tools command-line options* on page 2-18 in *Introducing ARM Compilation Tools*.

3.1.8 `--asm`

This option instructs the compiler to write a listing to a file of the disassembly of the machine code generated by the compiler.

Object code is generated when this option is selected. The link step is also performed, unless the `-c` option is chosen.

Note

To produce a disassembly of the machine code generated by the compiler, without generating object code, select `-S` instead of --asm.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Usage

The action of `--asm`, and the full name of the disassembly file produced, depends on the combination of options used:

Table 3-1 Compiling with the `--asm` option

Compiler option	Action
`--asm`	Writes a listing to a file of the disassembly of the compiled source. The link step is also performed, unless the `-c` option is used. The disassembly is written to a text file whose name defaults to the name of the input file with the filename extension `.s`.
`--asm -c`	As for `--asm`, except that the link step is not performed.
`--asm --interleave`	As for `--asm`, except that the source code is interleaved with the disassembly. The disassembly is written to a text file whose name defaults to the name of the input file with the filename extension `.txt`.
`--asm --multifile`	As for `--asm`, except that the compiler produces empty object files for the files merged into the main file.
`--asm -o` *`filename`*	As for `--asm`, except that the object file is named *`filename`*. The disassembly is written to the file *`filename.s`*. The name of the object file must not have the filename extension `.s`. If the filename extension of the object file is `.s`, the disassembly is written over the top of the object file. This might lead to unpredictable results.

See also

- *-c* on page 3-17
- *--interleave* on page 3-54
- *--multifile, --no_multifile* on page 3-67
- *-o filename* on page 3-69
- *-S* on page 3-85
- *Filename suffixes recognized by the compiler* on page 3-14 in *Using the Compiler*.

3.1.9 `--asm_dir=`*`directory_name`*

This option enables you to specify a directory for output assembler files.

Example

```
armcc -c --output_dir=obj -S f1.c f2.c --asm_dir=asm
```

Result:

```
asm/f1.s
asm/f2.s
obj/f1.o
obj/f2.o
```

See also

- *--depend_dir=directory_name* on page 3-27
- *--list_dir=directory_name* on page 3-61

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *--output_dir=directory_name* on page 3-73.

3.1.10 `--autoinline`, `--no_autoinline`

These options enable and disable automatic inlining of functions.

The compiler automatically inlines functions at the higher optimization levels where it is sensible to do so. The `-Ospace` and `-Otime` options, together with some other factors such as function size, influence how the compiler automatically inlines functions.

Selecting `-Otime`, in combination with various other factors, increases the likelihood that functions are inlined.

In general, when automatic inlining is enabled, the compiler inlines any function that is sensible to inline. When automatic inlining is disabled, only functions marked as `__inline` are candidates for inlining.

Usage

Use these options to control the automatic inlining of functions at the highest optimization levels (`-O2` and `-O3`).

Default

For optimization levels `-O0` and `-O1`, the default is `--no_autoinline`.

For optimization levels `-O2` and `-O3`, the default is `--autoinline`.

See also

- *--forceinline* on page 3-40
- *--inline, --no_inline* on page 3-53
- *-Onum* on page 3-71
- *-Ospace* on page 3-72
- *-Otime* on page 3-73
- *Default compiler options that are affected by optimization level* on page 4-41.

3.1.11 `--bigend`

This option instructs the compiler to generate code for an ARM processor using big-endian memory.

The ARM architecture defines the following big-endian modes:

BE8 Byte Invariant Addressing mode (ARMv6 and later).

BE32 Legacy big-endian mode.

The selection of BE8 versus BE32 is specified at link time.

Default

The compiler assumes `--littleend` unless `--bigend` is explicitly specified.

See also

- *--littleend* on page 3-61
- *--be8* on page 2-11 in the *Linker Reference*
- *--be32* on page 2-12 in the *Linker Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.12 `--bitband`

This option bit-bands all non **`const`** global structure objects. It enables a word of memory to be mapped to a single bit in the bit-band region. This enables efficient atomic access to single-bit values in SRAM and Peripheral regions of the memory architecture.

For peripherals that are width sensitive, byte, halfword, and word stores or loads to the alias space are generated for **`char`**, **`short`**, and **`int`** types of bitfields of bit-banded structs respectively.

Restrictions

The following restrictions apply:

- This option only affects **`struct`** types. Any union type or other aggregate type with a union as a member cannot be bit-banded.
- Members of structs cannot be bit-banded individually.
- Bit-banded accesses are generated only for single-bit bitfields.
- Bit-banded accesses are not generated for **`const`** objects, pointers, and local objects.
- Bit-banding is only available on some processors. For example, the Cortex-M4 and Cortex-M3 processors.

Example

In Example 3-1 the writes to bitfields `i` and `k` are bit-banded when compiled using the `--bitband` command-line option.

Example 3-1 Bit-banding example

```
typedef struct {
  int i : 1;
  int j : 2;
  int k : 1;
} BB;

BB value;

void update_value(void)
{
  value.i = 1;
  value.k = 1;
}
```

See also

- *__attribute__((bitband)) type attribute* on page 5-36
- *Compiler and processor support for bit-banding* on page 4-21 in *Using the Compiler*

- the *Technical Reference Manual* for your processor.

3.1.13 `--brief_diagnostics`, `--no_brief_diagnostics`

This option enables or disables the output of brief diagnostic messages by the compiler.

When enabled, the original source line is not displayed, and error message text is not wrapped if it is too long to fit on a single line.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Default

The default is `--no_brief_diagnostics`.

Example

```
/* main.c */
#include <stdio.h>
int main(void)
{
   printf(""Hello, world\n"); // Intentional quotation mark error
   return 0;
}
```

Compiling this code with `--brief_diagnostics` produces:

```
"main.c", line 5: Error:  #18: expected a ")"
"main.c", line 5: Error:  #7: unrecognized token
"main.c", line 5: Error:  #8: missing closing quote
"main.c", line 6: Error:  #65: expected a ";"
```

See also

- *--diag_error=tag[,tag,...]* on page 3-30
- *--diag_remark=tag[,tag,...]* on page 3-31
- *--diag_style={arm|ide|gnu}* on page 3-31
- *--diag_suppress=tag[,tag,...]* on page 3-32
- *--diag_warning=tag[,tag,...]* on page 3-33
- *--errors=filename* on page 3-36
- *--remarks* on page 3-82
- *-W* on page 3-96
- *--wrap_diagnostics, --no_wrap_diagnostics* on page 3-98
- Chapter 6 *Compiler Diagnostic Messages* in *Using the Compiler*.

3.1.14 `--bss_threshold=`*`num`*

This option controls the placement of small global ZI data items in sections. A *small global ZI data item* is an uninitialized data item that is eight bytes or less in size.

Syntax

`--bss_threshold=`*`num`*

Where:

`num` is either:

`0`	place small global ZI data items in ZI data sections
`8`	place small global ZI data items in RW data sections.

Usage

In ARM Compiler 4.1 and later, the compiler might place small global ZI data items in RW data sections as an optimization. In RVCT 2.0.1 and earlier, small global ZI data items were placed in ZI data sections by default.

Use `--bss_threshold=0` to emulate the behavior of RVCT 2.0.1 and earlier with respect to the placement of small global ZI data items in ZI data sections.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Note

Selecting the option `--bss_threshold=0` instructs the compiler to place *all* small global ZI data items in the current compilation module in a ZI data section. To place specific variables in:

- a ZI data section, use `__attribute__((zero_init))`
- a specific ZI data section, use a combination of `__attribute__((section("name")))` and `__attribute__((zero_init))`.

Default

If you do not specify a `--bss_threshold` option, the compiler assumes `--bss_threshold=8`.

Example

```
int glob1;          /* ZI (.bss) in RVCT 2.0.1 and earlier */
                    /* RW (.data) in RVCT 2.1 and later */
```

Compiling this code with `--bss_threshold=0` places `glob1` in a ZI data section.

See also

- *#pragma arm section [section_type_list]* on page 5-48
- *__attribute__((section("name"))) variable attribute* on page 5-42
- *__attribute__((zero_init)) variable attribute* on page 5-46.

3.1.15 -c

This option instructs the compiler to perform the compilation step, but not the link step.

Note

This option is different from the uppercase `-C` option.

Usage

The use of the `-c` option is recommended in projects with more than one source file.

See also

- *--asm* on page 3-12
- *--list* on page 3-59
- *-o filename* on page 3-69
- *-S* on page 3-85.

3.1.16 -C

This option instructs the compiler to retain comments in preprocessor output.

Choosing this option implicitly selects the option `-E`.

Note

This option is different from the lowercase `-c` option.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *-E* on page 3-35.

3.1.17 `--c90`

This option enables the compilation of C90 source code.

Usage

This option can also be combined with other source language command-line options. For example, `armcc --c90`.

Default

This option is implicitly selected for files having a suffix of `.c`, `.ac`, or `.tc`.

Note

If you are migrating from RVCT, be aware that filename extensions `.ac` and `.tc` are deprecated in ARM Compiler 4.1 and later.

See also

- *--c99*
- *--strict, --no_strict* on page 3-88
- *Source language modes* on page 2-3
- *Filename suffixes recognized by the compiler* on page 3-14 in *Using the Compiler*.

3.1.18 `--c99`

This option enables the compilation of C99 source code.

Usage

This option can also be combined with other source language command-line options. For example, `armcc --c99`.

Default

For files having a suffix of `.c`, `.ac`, or `.tc`, `--c90` applies by default.

See also

- *--c90*
- *--strict, --no_strict* on page 3-88
- *Source language modes* on page 2-3.

3.1.19 `--code_gen, --no_code_gen`

This option enables or disables the generation of object code.

When generation of object code is disabled, the compiler performs syntax-checking only, without creating an object file.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Default

The default is `--code_gen`.

3.1.20 `--compatible=name`

This option generates code that is compatible with multiple target architectures or processors.

Syntax

`--compatible=name`

Where:

name is the name of a target processor or architecture, or `NONE`. Processor and architecture names are not case-sensitive.

If multiple instances of this option are present on the command line, the last one specified overrides the previous instances.

Specify `--compatible=NONE` at the end of the command line to turn off all other instances of the option.

Usage

Using this option avoids the need to recompile the same source code for different targets. You could apply this use to a possible target upgrade where a different architecture or processor is to be used in the future, without having to separately recompile for that target.

See Table 3-2. The valid combinations are:

- `--cpu=CPU_from_group1 --compatible=CPU_from_group2`
- `--cpu=CPU_from_group2 --compatible=CPU_from_group1`.

Table 3-2 Compatible processor or architecture combinations

Group 1	`ARM7TDMI`, `4T`
Group 2	`Cortex-M0`, `Cortex-M1`, `Cortex-M3`, `Cortex-M4`, `7-M`, `6-M`, `6S-M`

No other combinations are permitted.

The effect is to compile code that is compatible with both `--cpu` and `--compatible`. This means that only Thumb1 instructions are used. (This is the intersection of the capabilities of group 1 and group 2.)

Note

Although the generated code is compatible with multiple targets, this code might be less efficient than compiling for a single target processor or architecture.

Example

This example gives code that is compatible with both the ARM7TDMI processor and the Cortex-M4 processor.

```
armcc --cpu=arm7tdmi --compatible=cortex-m4 myprog.c
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--cpu=name*.

3.1.21 `--compile_all_input, --no_compile_all_input`

These options enable and disable the suppression of filename extension processing, enabling the compiler to compile files with any filename extensions.

When enabled, the compiler suppresses filename extension processing entirely, treating all input files as if they have the suffix `.c`.

Default

The default is `--no_compile_all_input`.

See also

- *--link_all_input, --no_link_all_input* on page 3-58
- *Filename suffixes recognized by the compiler* on page 3-14 in *Using the Compiler*.

3.1.22 `--cpp`

This option enables the compilation of C++ source code.

Usage

This option can also be combined with other source language command-line options. For example, `armcc --cpp`.

Default

This option is implicitly selected for files having a suffix of `.cpp`, `.cxx`, `.c++`, `.cc`, or `.CC`.

See also

- *--anachronisms, --no_anachronisms* on page 3-7
- *--c90* on page 3-18
- *--c99* on page 3-18
- *--strict, --no_strict* on page 3-88
- *Source language modes* on page 2-3.

3.1.23 `--cpu=list`

This option lists the supported processor names that can be used with the `--cpu=name` option.

See also

- *--cpu=name*.

3.1.24 `--cpu=name`

This option enables code generation for the selected ARM processor.

Syntax

`--cpu=name`

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Where:

`name` is the name of a processor as shown on ARM data sheets, for example, `ARM7TDMI`.

Processor names are not case-sensitive.

Wildcard characters are not accepted.

Default

If you do not specify a `--cpu` option, the compiler assumes `--cpu=ARM7TDMI`.

To obtain a full list of CPU processors, use the `--cpu=list` option.

Usage

The following general points apply to processor options:

Processors

- Selecting the processor selects the appropriate architecture, *Floating-Point Unit* (FPU), and memory organization.
- The compiled code is optimized for the processor specified in the `--cpu` option. This enables the compiler to use specific coprocessors or instruction scheduling for optimum performance.

FPU

- Some specifications of `--cpu` imply an `--fpu` selection. For example, `--cpu=cortex-r4f` implies `--fpu=vfpv3_d16`.

  Note

  Any *explicit* FPU, set with `--fpu` on the command line, overrides an *implicit* FPU.

- If no `--fpu` option is specified and no `--cpu` option is specified, `--fpu=softvfp` is used.

ARM/Thumb

- Specifying a processor that supports Thumb instructions, such as `--cpu=ARM7TDMI`, does not make the compiler generate Thumb code. It only enables features of the processor to be used, such as long multiply. Use the `--thumb` option to generate Thumb code, unless the processor is a Thumb-only processor, for example Cortex-M4. In this case, `--thumb` is not required.

  Note

  Specifying the target processor might make the object code generated by the compiler incompatible with other ARM processors. For example, code compiled for architecture ARMv6 might not run on an ARM920T processor, if the compiled code includes instructions specific to ARMv6. Therefore, you must choose the lowest common denominator processor suited to your purpose.

- If you are compiling code that is intended for mixed ARM/Thumb systems for processors that support ARMv4T or ARMv5T, then you must specify the interworking option `--apcs=/interwork`. By default, this is enabled for processors that support ARMv5T or above.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- If you compile for Thumb, that is with the `--thumb` option on the command line, the compiler compiles as much of the code as possible using the Thumb instruction set. However, the compiler might generate ARM code for some parts of the compilation. For example, if you are compiling code for a Thumb-1 processor and using VFP, any function containing floating-point operations is compiled for ARM.

See also

- *--apcs=qualifer...qualifier* on page 3-7
- *--cpu=list* on page 3-20
- *--fpu=name* on page 3-44
- *--thumb* on page 3-90
- *__smc* on page 5-11
- *SMC* on page 3-141 in the *Assembler Reference*.

3.1.25 `--create_pch=filename`

This option instructs the compiler to create a *PreCompiled Header* (PCH) file with the specified filename.

This option takes precedence over all other PCH options.

Syntax

`--create_pch=filename`

Where:

`filename` is the name of the PCH file to be created.

See also

- *--pch* on page 3-75
- *--pch_dir=dir* on page 3-75
- *--pch_messages, --no_pch_messages* on page 3-76
- *--pch_verbose, --no_pch_verbose* on page 3-76
- *--use_pch=filename* on page 3-93
- *#pragma hdrstop* on page 5-52
- *#pragma no_pch* on page 5-54
- *PreCompiled Header (PCH) files* on page 4-29 in *Using the Compiler*.

3.1.26 `-Dname[(parm-list)][=def]`

This option defines the macro `name`.

Syntax

`-Dname[(parm-list)][=def]`

Where:

`name` Is the name of the macro to be defined.

`parm-list` Is an optional list of comma-separated macro parameters. By appending a macro parameter list to the macro name, you can define function-style macros.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

The parameter list must be enclosed in parentheses. When specifying multiple parameters, do not include spaces between commas and parameter names in the list.

Note

Parentheses might require escaping on UNIX systems.

`=def` Is an optional macro definition.

If `=def` is omitted, the compiler defines `name` as the value 1.

To include characters recognized as tokens on the command line, enclose the macro definition in double quotes.

Usage

Specifying `-Dname` has the same effect as placing the text `#define name` at the head of each source file.

Restrictions

The compiler defines and undefines macros in the following order:

1. compiler predefined macros
2. macros defined explicitly, using `-Dname`
3. macros explicitly undefined, using `-Uname`.

Example

Specifying the option:

```
-DMAX(X,Y)="((X > Y) ? X : Y)"
```

on the command line is equivalent to defining the macro:

```
#define MAX(X, Y) ((X > Y) ? X : Y)
```

at the head of each source file.

See also

- *-C* on page 3-17
- *-E* on page 3-35
- *-Uname* on page 3-91
- *Compiler predefines* on page 5-98.

3.1.27 `--data_reorder`, `--no_data_reorder`

This option enables or disables automatic reordering of top-level data items, for example global variables.

The compiler can save memory by eliminating wasted space between data items. However, `--data_reorder` can break legacy code, if the code makes invalid assumptions about ordering of data by the compiler.

The ISO C Standard does not guarantee data order, so you must try to avoid writing code that depends on any assumed ordering. If you require data ordering, place the data items into a structure.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Default

The default is optimization-level dependent:

`-O0`: `--no_data_reorder`

`-O1`, `-O2`, `-O3`: `--data_reorder`

See also

- *-Onum* on page 3-71
- *Default compiler options that are affected by optimization level* on page 4-41.

3.1.28 `--debug`, `--no_debug`

This option enables or disables the generation of debug tables for the current compilation.

The compiler produces the same code regardless of whether `--debug` is used. The only difference is the existence of debug tables.

Default

The default is `--no_debug`.

Using `--debug` does not affect optimization settings. By default, using the `--debug` option alone is equivalent to:

```
--debug --dwarf3 --debug_macros
```

See also

- *--debug_macros, --no_debug_macros*
- *--dwarf2* on page 3-35
- *--dwarf3* on page 3-35
- *-Onum* on page 3-71.

3.1.29 `--debug_macros`, `--no_debug_macros`

This option enables or disables the generation of debug table entries for preprocessor macro definitions.

Usage

Using `--no_debug_macros` might reduce the size of the debug image.

This option must be used with the `--debug` option.

Default

The default is `--debug_macros`.

See also

- *--debug, --no_debug*

3.1.30 `--default_definition_visibility=`*`visibility`*

This option controls the default ELF symbol visibility of **`extern`** variable and function definitions.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Syntax

`--default_definition_visibility=visibility`

Where:

visibility is `default`, `hidden`, `internal`, or `protected`.

Usage

Use `--default_definition_visibility=visibility` to force the compiler to use the specified ELF symbol visibility for all `extern` variables and functions defined in the source file, if they do not use `__declspec(dll*)`. Unlike `--hide_all`, `--no_hide_all`, this does not affect `extern` references.

Default

By default, `--default_definition_visibility=hidden`.

See also

- *--hide_all, --no_hide_all* on page 3-49

3.1.31 `--default_extension=ext`

This option enables you to change the filename extension for object files from the default extension (`.o`) to an extension of your choice.

Syntax

`--default_extension=ext`

Where:

ext is the filename extension of your choice.

Default

By default, the filename extension for object files is `.o`.

Example

The following example creates an object file called `test.obj`, instead of `test.o`:

```
armcc --default_extension=obj -c test.c
```

Note

The `-o filename` option overrides this. For example, the following command results in an object file named `test.o`:

```
armcc --default_extension=obj -o test.o -c test.c
```

3.1.32 `--dep_name, --no_dep_name`

This option enables or disables dependent name processing in C++.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

The C++ standard states that lookup of names in templates occurs:

- at the time the template is parsed, if the name is nondependent
- at the time the template is parsed, or at the time the template is instantiated, if the name is dependent.

When the option `--no_dep_name` is selected, the lookup of dependent names in templates can occur only at the time the template is instantiated. That is, the lookup of dependent names at the time the template is parsed is disabled.

Note

The option `--no_dep_name` is provided only as a migration aid for legacy source code that does not conform to the C++ standard. Its use is not recommended.

Mode

This option is effective only if the source language is C++.

Default

The default is `--dep_name`.

Restrictions

The option `--dep_name` cannot be combined with the option `--no_parse_templates`, because parsing is done by default when dependent name processing is enabled.

Errors

When the options `--dep_name` and `--no_parse_templates` are combined, the compiler generates an error.

See also

- *--parse_templates, --no_parse_templates* on page 3-74
- *Template instantiation* on page 6-12.

3.1.33 `--depend=filename`

This option instructs the compiler to write makefile dependency lines to a file during compilation.

Syntax

`--depend=filename`

Where:

`filename` is the name of the dependency file to be output.

Restrictions

If you specify multiple source files on the command line then any `--depend` option is ignored. No dependency file is generated in this case.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Usage

The output file is suitable for use by a make utility. To change the output format to be compatible with UNIX make utilities, use the `--depend_format` option.

See also

- *--depend_format=string*
- *--depend_system_headers, --no_depend_system_headers* on page 3-29
- *--depend_target=target* on page 3-30
- *--ignore_missing_headers* on page 3-50
- *--list* on page 3-59
- *-M* on page 3-64
- *--md* on page 3-65
- *--depend_single_line, --no_depend_single_line* on page 3-28
- *--phony_targets* on page 3-77

3.1.34 `--depend_dir=directory_name`

This option enables you to specify a directory for dependency output files.

Example

```
armcc -c --output_dir=obj f1.c f2.c --depend_dir=depend --depend=deps
```

Result:

```
depend/f1.d
depend/f2.d
obj/f1.o
obj/f2.o
```

See also

- *--asm_dir=directory_name* on page 3-13
- *--list_dir=directory_name* on page 3-61
- *--output_dir=directory_name* on page 3-73.

3.1.35 `--depend_format=string`

This option changes the format of output dependency files, for compatibility with some UNIX make programs.

Syntax

`--depend_format=`*`string`*

Where *`string`* is one of:

`unix`	generate dependency file entries using UNIX-style path separators.
`unix_escaped`	is the same as `unix`, but escapes spaces with \.
`unix_quoted`	is the same as `unix`, but surrounds path names with double quotes.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Usage

`unix` — On Windows systems, `--depend_format=unix` forces the use of UNIX-style path names. That is, the UNIX-style path separator symbol / is used in place of \.

On UNIX systems, `--depend_format=unix` has no effect.

`unix_escaped` — On Windows systems, `--depend_format=unix_escaped` forces unix-style path names, and escapes spaces with \.

On UNIX systems, `--depend_format=unix_escaped` with escapes spaces with \.

`unix_quoted` — On Windows systems, `--depend_format=unix_quoted` forces unix-style path names and surrounds them with "".

On UNIX systems, `--depend_format=unix_quoted` surrounds path names with `""`.

Default

If you do not specify a `--depend_format` option, then the format of output dependency files depends on your choice of operating system:

Windows — On Windows systems, the default is to use either Windows-style paths or UNIX-style paths, whichever is given.

UNIX — On UNIX systems, the default is `--depend_format=unix`.

Example

On a Windows system, compiling a file `main.c` containing the line:

```
#include "..\include\header files\common.h"
```

using the options `--depend=depend.txt --depend_format=unix_escaped` produces a dependency file `depend.txt` containing the entries:

```
main.axf: main.c
main.axf: ../include/header\ files/common.h
```

See also

- *--depend=filename* on page 3-26
- *--depend_system_headers, --no_depend_system_headers* on page 3-29
- *--depend_target=target* on page 3-30
- *--ignore_missing_headers* on page 3-50
- *-M* on page 3-64
- *--md* on page 3-65
- *--phony_targets* on page 3-77

3.1.36 `--depend_single_line, --no_depend_single_line`

This option determines the format of the makefile dependency lines output by the compiler. `--depend_single_line` instructs the compiler to format the makefile with one dependency line for each compilation unit. The compiler wraps long lines to improve readability. `--no_depend_single_line` instructs the compiler to format the makefile with one line for each include file or source file.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Default

The default is `--no_depend_single_line`.

Example

```
/* hello.c */
#include <stdio.h>
int main(void)
{
    printf("Hello, world!\n");
    return 0;
}
```

Compiling this code with `armcc hello.c -M --depend_single_line` produces:

```
__image.axf: hello.c_...\include\...\stdio.h
```

Compiling this code with `armcc hello.c -M --no_depend_single_line` produces:

```
__image.axf: hello.c
__image.axf: ...\include\...\stdio.h
```

See also

- *--depend=filename* on page 3-26
- *--depend_format=string* on page 3-27
- *--depend_target=target* on page 3-30
- *--ignore_missing_headers* on page 3-50
- *-M* on page 3-64
- *--md* on page 3-65
- *--phony_targets* on page 3-77.

3.1.37 `--depend_system_headers`, `--no_depend_system_headers`

This option enables or disables the output of system include dependency lines when generating makefile dependency information using either the `-M` option or the `--md` option.

Default

The default is `--depend_system_headers`.

Example

```
/* hello.c */
#include <stdio.h>
int main(void)
{
    printf("Hello, world!\n");
    return 0;
}
```

Compiling this code with the option `-M` produces:

```
__image.axf: hello.c
__image.axf: ...\include\...\stdio.h
```

Compiling this code with the options `-M --no_depend_system_headers` produces:

```
__image.axf: hello.c
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--depend=filename* on page 3-26
- *--depend_format=string* on page 3-27
- *--depend_target=target*
- *--ignore_missing_headers* on page 3-50
- *-M* on page 3-64
- *--md* on page 3-65
- *--phony_targets* on page 3-77

3.1.38 `--depend_target=`*`target`*

This option sets the target for makefile dependency generation.

Usage

Use this option to override the default target.

Restriction

This option is analogous to `-MT` in GCC. However, behavior differs when specifying multiple targets. For example, `gcc -M -MT target1 -MT target2 file.c` might give a result of `target1 target2: file.c header.h`, whereas `--depend_target=target1 --depend_target=target2` treats `target2` as the target.

See also

- *--depend=filename* on page 3-26
- *--depend_format=string* on page 3-27
- *--depend_system_headers, --no_depend_system_headers* on page 3-29
- *--ignore_missing_headers* on page 3-50
- *-M* on page 3-64
- *--md* on page 3-65
- *--phony_targets* on page 3-77

3.1.39 `--diag_error=`*`tag`*`[,`*`tag`*`,...]`

This option sets diagnostic messages that have a specific tag to error severity.

Note

This option has the `#pragma` equivalent `#pragma diag_error`.

Syntax

`--diag_error=`*`tag`*`[,`*`tag`*`,...]`

Where *`tag`* can be:

- a diagnostic message number to set to error severity
- `warning`, to treat all warnings as errors.

Usage

The severity of the following types of diagnostic messages can be changed:

- Messages with the number format `#`*`nnnn`*`-D`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- Warning messages with the number format `CnnnnW`.

See also

- *--diag_remark=tag[,tag,...]*
- *--diag_suppress=tag[,tag,...]* on page 3-32
- *--diag_warning=tag[,tag,...]* on page 3-33
- *#pragma diag_error tag[,tag,...]* on page 5-50
- *Options that change the severity of compiler diagnostic messages* on page 6-4 in *Using the Compiler*.

3.1.40 `--diag_remark=tag[,tag,...]`

This option sets the diagnostic messages that have the specified tags to Remark severity.

The `--diag_remark` option behaves analogously to `--diag_errors`, except that the compiler sets the diagnostic messages having the specified tags to Remark severity rather than Error severity.

Note

Remarks are not displayed by default. To see remark messages, use the compiler option `--remarks`.

Note

This option has the `#pragma` equivalent `#pragma diag_remark`.

Syntax

```
--diag_remark=tag[,tag,...]
```

Where:

`tag[,tag,...]` is a comma-separated list of diagnostic message numbers specifying the messages whose severities are to be changed.

See also

- *--diag_error=tag[,tag,...]* on page 3-30
- *--diag_suppress=tag[,tag,...]* on page 3-32
- *--diag_warning=tag[,tag,...]* on page 3-33
- *--remarks* on page 3-82
- *#pragma diag_remark tag[,tag,...]* on page 5-50
- *Options that change the severity of compiler diagnostic messages* on page 6-4 in *Using the Compiler*.

3.1.41 `--diag_style={arm|ide|gnu}`

This option specifies the style used to display diagnostic messages.

Syntax

```
--diag_style=string
```

Where *string* is one of:

`arm` Display messages using the ARM compiler style.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

`ide` Include the line number and character count for any line that is in error. These values are displayed in parentheses.

`gnu` Display messages in the format used by `gcc`.

Default

If you do not specify a `--diag_style` option, the compiler assumes `--diag_style=arm`.

Usage

Choosing the option `--diag_style=ide` implicitly selects the option `--brief_diagnostics`. Explicitly selecting `--no_brief_diagnostics` on the command line overrides the selection of `--brief_diagnostics` implied by `--diag_style=ide`.

Selecting either the option `--diag_style=arm` or the option `--diag_style=gnu` does not imply any selection of `--brief_diagnostics`.

See also

- *--diag_error=tag[,tag,...]* on page 3-30
- *--diag_remark=tag[,tag,...]* on page 3-31
- *--diag_suppress=tag[,tag,...]*
- *--diag_warning=tag[,tag,...]* on page 3-33
- *Options that change the severity of compiler diagnostic messages* on page 6-4 in *Using the Compiler*.

3.1.42 `--diag_suppress=tag[,tag,...]`

This option disables diagnostic messages that have the specified tags.

The `--diag_suppress` option behaves analogously to `--diag_errors`, except that the compiler suppresses the diagnostic messages having the specified tags rather than setting them to have error severity.

Note

This option has the `#pragma` equivalent `#pragma diag_suppress`.

Syntax

`--diag_suppress=tag[,tag,...]`

Where *tag* can be:

- a diagnostic message number to be suppressed
- `error`, to suppress all downgradeable errors
- `warning`, to suppress all warnings.

See also

- *--diag_error=tag[,tag,...]* on page 3-30
- *--diag_remark=tag[,tag,...]* on page 3-31
- *--diag_warning=tag[,tag,...]* on page 3-33
- *#pragma diag_suppress tag[,tag,...]* on page 5-51
- *Compiler diagnostics* on page 6-2 in *Using the Compiler*
- *Prefix letters in compiler diagnostic messages* on page 6-5 in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.43 `--diag_suppress=optimizations`

This option suppresses diagnostic messages for high-level optimizations.

Default

By default, optimization messages have Remark severity. Specifying `--diag_suppress=optimizations` suppresses optimization messages.

Note

Use the `--remarks` option to see optimization messages having Remark severity.

Usage

The compiler performs certain high-level vector and scalar optimizations when compiling at the optimization level `-O3 -Otime`, for example, loop unrolling. Use this option to suppress diagnostic messages relating to these high-level optimizations.

Example

```
int factorial(int n)
{
    int result=1;
    while (n > 0)
        result *= n--;
    return result;
}
```

Compiling this code with the options `-O3 -Otime --remarks --diag_suppress=optimizations` suppresses optimization messages.

See also

- *--diag_suppress=tag[,tag,...]* on page 3-32
- *--diag_warning=optimizations* on page 3-34
- *-Onum* on page 3-71
- *-Otime* on page 3-73
- *--remarks* on page 3-82.

3.1.44 `--diag_warning=tag[,tag,...]`

This option sets diagnostic messages that have the specified tags to warning severity.

The `--diag_warning` option behaves analogously to `--diag_errors`, except that the compiler sets the diagnostic messages having the specified tags to warning severity rather than error severity.

Note

This option has the `#pragma` equivalent `#pragma diag_warning`.

Syntax

`--diag_warning=tag[,tag,...]`

Where *tag* can be:

- a diagnostic message number to set to warning severity

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- `error`, to downgrade the severity of all downgradeable errors to warnings.

Example

`--diag_warning=A1234,error` causes message `A1234` and all downgradeable errors to be treated as warnings, providing changing the severity of `A1234` is permitted.

See also

- *--diag_error=tag[,tag,...]* on page 3-30
- *--diag_remark=tag[,tag,...]* on page 3-31
- *--diag_suppress=tag[,tag,...]* on page 3-32
- *#pragma diag_warning tag[, tag, ...]* on page 5-51
- *Options that change the severity of compiler diagnostic messages* on page 6-4 in *Using the Compiler*.

3.1.45 `--diag_warning=optimizations`

This option sets high-level optimization diagnostic messages to have Warning severity.

Default

By default, optimization messages have Remark severity.

Usage

The compiler performs certain high-level vector and scalar optimizations when compiling at the optimization level `-O3 -Otime`, for example, loop unrolling. Use this option to display diagnostic messages relating to these high-level optimizations.

Example

```
int factorial(int n)
{
    int result=1;
    while (n > 0)
        result *= n--;
    return result;
}
```

Compiling this code with the options `--cpu=Cortex-R4F -O3 -Otime --diag_warning=optimizations` generates optimization warning messages.

See also

- *--diag_suppress=optimizations* on page 3-33
- *--diag_warning=tag[,tag,...]* on page 3-33
- *-Onum* on page 3-71
- *-Otime* on page 3-73.

3.1.46 `--dollar`, `--no_dollar`

This option instructs the compiler to accept or reject dollar signs, `$`, in identifiers.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Default

If the options `--strict` or `--strict_warnings` are specified, the default is `--no_dollar`. Otherwise, the default is `--dollar`.

See also

- *Dollar signs in identifiers* on page 4-11
- *--strict, --no_strict* on page 3-88.

3.1.47 `--dwarf2`

This option instructs the compiler to use DWARF 2 debug table format.

Default

The compiler assumes `--dwarf3` unless `--dwarf2` is explicitly specified.

See also

- *--dwarf3*.

3.1.48 `--dwarf3`

This option instructs the compiler to use DWARF 3 debug table format.

Default

The compiler assumes `--dwarf3` unless `--dwarf2` is explicitly specified.

See also

- *--dwarf2*.

3.1.49 `-E`

This option instructs the compiler to execute only the preprocessor step.

By default, output from the preprocessor is sent to the standard output stream and can be redirected to a file using standard UNIX and MS-DOS notation.

You can also use the `-o` option to specify a file for the preprocessed output. By default, comments are stripped from the output. The preprocessor accepts source files with any extension, for example, `.o`, `.s`, and `.txt`.

To generate interleaved macro definitions and preprocessor output, use `-E --list_macros`.

Example

```
armcc -E source.c > raw.c
```

See also

- *-C* on page 3-17
- *--list_macros* on page 3-61
- *--md* on page 3-65
- *-o filename* on page 3-69
- *--old_style_preprocessing* on page 3-72

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *-P* on page 3-74.

3.1.50 `--emit_frame_directives`, `--no_emit_frame_directives`

This option instructs the compiler to place DWARF `FRAME` directives into disassembly output.

Default

The default is `--no_emit_frame_directives`.

Examples

```
armcc --asm --emit_frame_directives foo.c
```

```
armcc -S emit_frame_directives foo.c
```

See also

- *--asm* on page 3-12
- *-S* on page 3-85
- *Frame directives* on page 5-37 in *Using the Assembler*.

3.1.51 `--enum_is_int`

This option forces the size of all enumeration types to be at least four bytes.

Note

The `--enum_is_int` option is not recommended for general use.

Default

This option is switched off by default. The smallest data type that can hold the values of all enumerators is used.

See also

- *--interface_enums_are_32_bit* on page 3-53
- *Structures, unions, enumerations, and bitfields* on page 6-6.

3.1.52 `--errors=`*`filename`*

This option redirects the output of diagnostic messages from `stderr` to the specified errors file.

Syntax

`--errors=`*`filename`*

Where:

`filename` is the name of the file to which errors are to be redirected.

Diagnostics that relate to problems with the command options are not redirected, for example, if you type an option name incorrectly. However, if you specify an invalid argument to an option, for example `--cpu=999`, the related diagnostic is redirected to the specified *`filename`*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Usage

This option is useful on systems where output redirection of files is not well supported.

See also

- *--brief_diagnostics, --no_brief_diagnostics* on page 3-15
- *--diag_error=tag[,tag,...]* on page 3-30
- *--diag_remark=tag[,tag,...]* on page 3-31
- *--diag_style={arm|ide|gnu}* on page 3-31
- *--diag_suppress=tag[,tag,...]* on page 3-32
- *--diag_warning=tag[,tag,...]* on page 3-33
- *--remarks* on page 3-82
- *-W* on page 3-96
- *--wrap_diagnostics, --no_wrap_diagnostics* on page 3-98
- Chapter 6 *Compiler Diagnostic Messages* in *Using the Compiler*.

3.1.53 `--exceptions`, `--no_exceptions`

This option enables or disables exception handling.

In C++, the `--exceptions` option enables the use of throw and try/catch, causes function exception specifications to be respected, and causes the compiler to emit unwinding tables to support exception propagation at runtime.

In C++, when the `--no_exceptions` option is specified, throw and try/catch are not permitted in source code. However, function exception specifications are still parsed, but most of their meaning is ignored.

In C, the behavior of code compiled with `--no_exceptions` is undefined if an exception is thrown through the compiled functions. You must use `--exceptions`, if you want exceptions to propagate correctly though C functions.

Default

The default is `--no_exceptions`.

See also

- *--exceptions_unwind, --no_exceptions_unwind*.

3.1.54 `--exceptions_unwind`, `--no_exceptions_unwind`

This option enables or disables function unwinding for exception-aware code. This option is only effective if `--exceptions` is enabled.

When you use `--no_exceptions_unwind` and `--exceptions` then no exception can propagate through the compiled functions. `std::terminate` is called instead.

Default

The default is `--exceptions_unwind`.

See also

- *--exceptions, --no_exceptions*
- *Function unwinding at runtime* on page 6-15.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.55 `--export_all_vtbl`, `--no_export_all_vtbl`

This option controls how dynamic symbols are exported in C++.

Mode

This option is effective only if the source language is C++.

Default

The default is `--no_export_all_vtbl`.

Usage

Use the option `--export_all_vtbl` to export all virtual function tables and RTTI for classes with a key function. A *key function* is the first virtual function of a class, in declaration order, that is not inline, and is not pure virtual.

Note

You can disable export for specific classes by using `__declspec(notshared)`.

See also

- *__declspec(notshared)* on page 5-23
- *--import_all_vtbl* on page 3-52.

3.1.56 `--export_defs_implicitly`, `--no_export_defs_implicitly`

This option controls how dynamic symbols are exported.

Default

The default is `--no_export_defs_implicitly`.

Usage

Use the option `--export_defs_implicitly` to export definitions where the prototype is marked `__declspec(dllimport)`.

See also

- *__declspec(dllimport)* on page 5-21.

3.1.57 `--extended_initializers`, `--no_extended_initializers`

These options enable and disable the use of extended constant initializers even when compiling with `--strict` or `--strict_warnings`.

When certain nonportable but widely supported constant initializers such as the cast of an address to an integral type are used, `--extended_initializers` causes the compiler to produce the same general warning concerning constant initializers that it normally produces in nonstrict mode, rather than specific errors stating that the expression must have a constant value or have arithmetic type.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Default

The default is `--no_extended_initializers` when compiling with `--strict` or `--strict_warnings`.

The default is `--extended_initializers` when compiling in nonstrict mode.

See also

- *--strict, --no_strict* on page 3-88
- *--strict_warnings* on page 3-89
- *Constant expressions* on page 4-9.

3.1.58 `--feedback=filename`

This option enables the efficient elimination of unused functions, and on the ARMv4T architecture, enables reduction of compilation required for interworking.

Syntax

`--feedback=filename`

Where:

filename is the feedback file created by a previous execution of the ARM linker.

Usage

You can perform multiple compilations using the same feedback file. The compiler places each unused function identified in the feedback file into its own ELF section in the corresponding object file.

The feedback file contains information about a previous build. Because of this:

- The feedback file might be out of date. That is, a function previously identified as being unused might be used in the current source code. The linker removes the code for an unused function only if it is not used in the current source code.

  Note

  - For this reason, eliminating unused functions using linker feedback is a safe optimization, but there might be a small impact on code size.
  - The usage requirements for reducing compilation required for interworking are more strict than for eliminating unused functions. If you are reducing interworking compilation, it is critical that you keep your feedback file up to date with the source code that it was generated from.

- You have to do a full compile and link at least twice to get the maximum benefit from linker feedback. However, a single compile and link using feedback from a previous build is usually sufficient.

See also

- *--split_sections* on page 3-87
- *--feedback_type=type* on page 2-50 in the *Linker Reference*
- *Linker feedback during compilation* on page 3-22 in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.59 --force_new_nothrow, --no_force_new_nothrow

This option controls the behavior of new expressions in C++.

The C++ standard states that only a no throw operator new declared with throw() is permitted to return NULL on failure. Any other operator new is never permitted to return NULL and the default operator new throws an exception on failure.

If you use --force_new_nothrow, the compiler treats expressions such as new T(...args...), that use the global ::operator new or ::operator new[], as if they are new (std::nothrow) T(...args...).

--force_new_nothrow also causes any class-specific operator new or any overloaded global operator new to be treated as no throw.

Note

The option --force_new_nothrow is provided only as a migration aid for legacy source code that does not conform to the C++ standard. Its use is not recommended.

Mode

This option is effective only if the source language is C++.

Default

The default is --no_force_new_nothrow.

Example

```
struct S
{
    void* operator new(std::size_t);
    void* operator new[](std::size_t);
};
void *operator new(std::size_t, int);
```

With the --force_new_nothrow option in effect, this is treated as:

```
struct S
{
    void* operator new(std::size_t) throw();
    void* operator new[](std::size_t) throw();
};
void *operator new(std::size_t, int) throw();
```

See also

- *Using the ::operator new function* on page 6-11.

3.1.60 --forceinline

This option forces all inline functions to be treated as if they are qualified with __forceinline.

Inline functions are functions that are qualified with inline or __inline. In C++, inline functions are functions that are defined inside a struct, class, or union definition.

If you use --forceinline, the compiler always attempts to inline those functions, if possible. However, it does not inline a function if doing so causes problems. For example, a recursive function is inlined into itself only once.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

`__forceinline` behaves like `__inline` except that the compiler tries harder to do the inlining.

See also

- *--autoinline, --no_autoinline* on page 3-14
- *--inline, --no_inline* on page 3-53
- *__forceinline* on page 5-5
- *__inline* on page 5-7
- *Inline functions on page 5-29* in *Using the Compiler*.

3.1.61 `--fp16_format=format`

This option enables the use of half-precision floating-point numbers as an optional extension to the VFPv3 architecture. If a format is not specified, use of the `__fp16` data type is faulted by the compiler.

Syntax

`--fp16_format=format`

Where *format* is one of:

`alternative`	An alternative to `ieee` that provides additional range, but has no NaN or infinity values.
`ieee`	Half-precision binary floating-point format defined by IEEE 754r, a revision to the IEEE 754 standard.
`none`	This is the default setting. It is equivalent to not specifying a format and means that the compiler faults use of the `__fp16` data type.

Restrictions

The following restrictions apply when you use the `__fp16` data type:

- When used in a C or C++ expression, an `__fp16` type is promoted to single precision. Subsequent promotion to double precision can occur if required by one of the operands.
- A single precision value can be converted to `__fp16`. A double precision value is converted to single precision and then to `__fp16`, that could involve double rounding. This reflects the lack of direct double-to-16-bit conversion in the ARM architecture.
- When using `fpmode=fast`, no floating-point exceptions are raised when converting to and from half-precision floating-point format.
- Function formal arguments cannot be of type `__fp16`. However, pointers to variables of type `__fp16` can be used as function formal argument types.
- `__fp16` values can be passed as actual function arguments. In this case, they are converted to single-precision values.
- `__fp16` cannot be specified as the return type of a function. However, a pointer to an `__fp16` type can be used as a return type.
- An `__fp16` value is converted to a single-precision or double-precision value when used as a return value for a function that returns a `float` or `double`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--fpmode=model*
- *Compiler and library support for half-precision floating-point numbers on page 5-61* of *Using the Compiler*.

3.1.62 `--fpmode=model`

This option specifies the floating-point conformance, and sets library attributes and floating-point optimizations.

Syntax

`--fpmode=model`

Where `model` is one of:

`ieee_full` All facilities, operations, and representations guaranteed by the IEEE standard are available in single and double-precision. Modes of operation can be selected dynamically at runtime.

This defines the symbols:

```
__FP_IEEE
__FP_FENV_EXCEPTIONS
__FP_FENV_ROUNDING
__FP_INEXACT_EXCEPTION
```

`ieee_fixed`

IEEE standard with round-to-nearest and no inexact exceptions.

This defines the symbols:

```
__FP_IEEE
__FP_FENV_EXCEPTIONS
```

`ieee_no_fenv`

IEEE standard with round-to-nearest and no exceptions. This mode is stateless and is compatible with the Java floating-point arithmetic model.

This defines the symbol `__FP_IEEE`.

`none` The compiler permits `--fpmode=none` as an alternative to `--fpu=none`, indicating that source code is not permitted to use floating-point types of any kind.

`std` IEEE finite values with denormals flushed to zero, round-to-nearest, and no exceptions. This is compatible with standard C and C++ and is the default option.

Normal finite values are as predicted by the IEEE standard. However:

- NaNs and infinities might not be produced in all circumstances defined by the IEEE model. When they are produced, they might not have the same sign.
- The sign of zero might not be that predicted by the IEEE model.

`fast` Perform more aggressive floating-point optimizations that might cause a small loss of accuracy to provide a significant performance increase. This option defines the symbol `__FP_FAST`.

This option results in behavior that is not fully compliant with the ISO C or C++ standard. However, numerically robust floating-point programs are expected to behave correctly.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A number of transformations might be performed, including:

- Double-precision math functions might be converted to single precision equivalents if all floating-point arguments can be exactly represented as single precision values, and the result is immediately converted to a single-precision value.

  This transformation is only performed when the selected library contains the single-precision equivalent functions, for example, when the selected library is `armcc`.

  For example:

  ```
  float f(float a)
  {
      return sqrt(a);
  }
  ```

  is transformed to

  ```
  float f(float a)
  {
      return sqrtf(a);
  }.
  ```

- Double-precision floating-point expressions that are narrowed to single-precision are evaluated in single-precision when it is beneficial to do so. For example, `float y = (float)(x + 1.0)` is evaluated as `float y = (float)x + 1.0f`.
- Division by a floating-point constant is replaced by multiplication with the inverse. For example, `x / 3.0` is evaluated as `x * (1.0 / 3.0)`.
- It is not guaranteed that the value of `errno` is compliant with the ISO C or C++ standard after math functions have been called. This enables the compiler to inline the VFP square root instructions in place of calls to `sqrt()` or `sqrtf()`.

Note

Initialization code might be required to enable the VFP. See *Limitations on hardware handling of floating-point arithmetic on page 5-59* in *Using the Compiler* for more information.

Default

By default, `--fpmode=std` applies.

See also

- *--fpu=name* on page 3-44
- *Using VFP with RVDS, Application Note 133*, http://infocenter/help/index.jsp?topic=/com.arm.doc.dai0133-

3.1.63 `--fpu=list`

This option lists the supported FPU architecture names that you can use with the `--fpu=name` option.

Deprecated options are not listed.

See also

- *--fpu=name* on page 3-44.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.64 `--fpu=name`

This option enables you to specify the target FPU architecture.

If you specify this option, it overrides any implicit FPU option that appears on the command line, for example, where you use the `--cpu` option.

To obtain a full list of FPU architectures use the `--fpu=list` option.

Syntax

`--fpu=name`

Where *name* is one of:

`none` Selects no floating-point option. No floating-point code is to be used. This produces an error if your code contains **`float`** types.

`vfpv` This is a synonym for `vfpv2`.

`vfpv2` Selects a hardware vector floating-point unit conforming to architecture VFPv2.

Note

If you enter `armcc --thumb --fpu=vfpv2` on the command line, the compiler compiles as much of the code using the Thumb instruction set as possible, but hard floating-point sensitive functions are compiled to ARM code. In this case, the value of the predefine `__thumb` is not correct.

`vfpv3` Selects a hardware vector floating-point unit conforming to architecture VFPv3. VFPv3 is backwards compatible with VFPv2 except that VFPv3 cannot trap floating-point exceptions.

`vfpv3_fp16` Selects a hardware vector floating-point unit conforming to architecture VFPv3 that also provides the half-precision extensions.

`vfpv3_d16` Selects a hardware vector floating-point unit conforming to VFPv3-D16 architecture.

`vfpv3_d16_fp16`

Selects a hardware vector floating-point unit conforming to VFPv3-D16 architecture, that also provides the half-precision extensions.

`vfpv4` Selects a hardware floating-point unit conforming to FPv4 architecture.

`vfpv4_d16` Selects a hardware floating-point unit conforming to the VFPv4-D16 architecture.

`fpv4-sp` Selects a hardware floating-point unit conforming to the single precision variant of the FPv4 architecture.

`softvfp` Selects software floating-point support where floating-point operations are performed by a floating-point library, `fplib`. This is the default if you do not specify a `--fpu` option, or if you select a CPU that does not have an FPU.

`softvfp+vfpv2`

Selects a hardware vector floating-point unit conforming to VFPv2, with software floating-point linkage. Select this option if you are interworking Thumb code with ARM code on a system that implements a VFP unit.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

If you select this option:

- Compiling with `--thumb` behaves in a similar way to `--fpu=softvfp` except that it links with floating-point libraries that use VFP instructions.
- Compiling with `--arm` behaves in a similar way to `--fpu=vfpv2` except that all functions are given software floating-point linkage. This means that functions pass and return floating-point arguments and results in the same way as `--fpu=softvfp`, but use VFP instructions internally.

Note

If you specify `softvfp+vfpv2` with the `--arm` or `--thumb` option for C code, it ensures that your interworking floating-point code is compiled to use software floating-point linkage.

`softvfp+vfpv3`

Selects a hardware vector floating-point unit conforming to VFPv3, with software floating-point linkage. Select this option if you are interworking Thumb code with ARM code on a system that implements a VFPv3 unit.

`softvfp+vfpv3_fp16`

Selects a hardware vector floating-point unit conforming to VFPv3-fp16, with software floating-point linkage.

`softvfp+vfpv3_d16`

Selects a hardware vector floating-point unit conforming to VFPv3-D16, with software floating-point linkage.

`softvfp+vfpv3_d16_fp16`

Selects a hardware vector floating-point unit conforming to `vfpv3_d16_fp16`, with software floating-point linkage.

`softvfp+vfpv4`

Selects a hardware floating-point unit conforming to FPv4, with software floating-point linkage.

`softvfp+vfpv4_d16`

Selects a hardware floating-point unit conforming to VFPv4-D16, with software floating-point linkage.

`softvfp+fpv4-sp`

Selects a hardware floating-point unit conforming to FPv4-SP, with software floating-point linkage.

Usage

Any FPU explicitly selected using the `--fpu` option always overrides any FPU implicitly selected using the `--cpu` option. For example, the option `--cpu=Cortex-R4F --fpu=softvfp` generates code that uses the software floating-point library `fplib`, even though the choice of CPU implies the use of architecture VFPv3-D16.

To control floating-point linkage without affecting the choice of FPU, you can use `--apcs=/softfp` or `--apcs=/hardfp`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Restrictions

The compiler only permits hardware VFP architectures (for example, `--fpu=vfpv3`, `--fpu=softvfp+vfpv2`), to be specified when `MRRC` and `MCRR` instructions are supported in the processor instruction set. `MRRC` and `MCRR` instructions are not supported in 4, 4T, 5T and 6-M. Therefore, the compiler does not allow the use of these CPU architectures with hardware VFP architectures.

Other than this, the compiler does not check that `--cpu` and `--fpu` combinations are valid. Beyond the scope of the compiler, additional architectural constraints apply. For example, VFPv3 is not supported with architectures prior to ARMv7. Therefore, the combination of `--fpu` and `--cpu` options permitted by the compiler does not necessarily translate to the actual device in use.

The compiler only generates scalar floating-point operations. If you want to use VFP vector operations, you must do this using assembly code.

Default

The default target FPU architecture is derived from the use of the `--cpu` option.

If the CPU specified with `--cpu` has a VFP coprocessor, the default target FPU architecture is the VFP architecture for that CPU. If a VFP coprocessor is present, VFP instructions are generated.

If there is no VFP coprocessor, the compiler generates code that makes calls to the software floating-point library `fplib` to carry out floating-point operations.

See also

- *--apcs=qualifer...qualifier* on page 3-7
- *--arm* on page 3-11
- *--cpu=name* on page 3-20
- *--fpmode=model* on page 3-42
- *--thumb* on page 3-90
- *__softfp* on page 5-12
- *Vector Floating-Point (VFP) architectures on page 5-58* in *Using the Compiler*
- *Compiler support for floating-point computations and linkage on page 5-63* in *Using the Compiler*
- *MRC, MRC2, MRRC and MRRC2* on page 3-127 in the *Assembler Reference*
- *MCR, MCR2, MCRR, and MCRR2* on page 3-126 in the *Assembler Reference*.

3.1.65 `--friend_injection`, `--no_friend_injection`

This option controls the visibility of `friend` declarations in C++.

In C++, it controls whether or not the name of a class or function that is declared only in `friend` declarations is visible when using the normal lookup mechanisms.

When `friend` names are declared, they are visible to these lookups. When `friend` names are not declared as required by the standard, function names are visible only when using argument-dependent lookup, and class names are never visible.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

——— **Note** ———

The option `--friend_injection` is provided only as a migration aid for legacy source code that does not conform to the C++ standard. Its use is not recommended.

Mode

This option is effective only if the source language is C++.

Default

The default is `--no_friend_injection`.

See also

- *friend* on page 4-13.

3.1.66 `-g`

This option enables the generation of debug tables for the current compilation.

The compiler produces the same code regardless of whether `-g` is used. The only difference is the existence of debug tables.

Using `-g` does not affect optimization settings.

Default

By default, using the `-g` option alone is equivalent to:

```
-g --dwarf3 --debug_macros
```

See also

- *--debug, --no_debug* on page 3-24
- *--debug_macros, --no_debug_macros* on page 3-24
- *--dwarf2* on page 3-35
- *--dwarf3* on page 3-35
- *-Onum* on page 3-71.

3.1.67 `--global_reg=`*`reg_name`*`[,`*`reg_name`*`,...]`

This option treats the specified register names as fixed registers, and prevents the compiler from using them in the code that is generated.

——— **Note** ———

Try to avoid using this option, because it restricts the compiler in terms of register allocation and can potentially give a negative effect on code generation and performance.

Syntax

```
--global_reg=reg_name[,reg_name,...]
```

Where *`reg_name`* is the AAPCS name of the register, denoted by an integer value in the range `1` to `8`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Register names 1 to 8 map sequentially onto registers r4 to r11.

If *reg_name* is unspecified, the compiler faults use of --global_reg.

Restrictions

This option has the same restrictions as the __global_reg storage class specifier.

Example

--global_reg=1,4,5 // reserve registers r4, r7 and r8 respectively

See also

- *__global_reg* on page 5-5
- *ARM Software Development Toolkit Reference Guide*.

3.1.68 --guiding_decls, --no_guiding_decls

This option enables or disables the recognition of guiding declarations for template functions in C++.

A *guiding declaration* is a function declaration that matches an instance of a function template but has no explicit definition because its definition derives from the function template.

If --no_guiding_decls is combined with --old_specializations, a specialization of a nonmember template function is not recognized. It is treated as a definition of an independent function.

Note

The option --guiding_decls is provided only as a migration aid for legacy source code that does not conform to the C++ standard. Its use is not recommended.

Mode

This option is effective only if the source language is C++.

Default

The default is --no_guiding_decls.

Example

```
template <class T> void f(T)
{
    ...
}
void f(int);
```

When regarded as a guiding declaration, f(int) is an instance of the template. Otherwise, it is an independent function so you must supply a definition.

See also

- *--apcs=qualifer...qualifier* on page 3-7
- *--old_specializations, --no_old_specializations* on page 3-72.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.69 `--help`

This option displays a summary of the main command-line options.

Default

`--help` applies by default if you fail to specify any command-line options or source files.

See also

- *--show_cmdline* on page 3-85.
- *--vsn* on page 3-96

3.1.70 `--hide_all`, `--no_hide_all`

This option enables you to control symbol visibility when building SVr4 shared objects.

Usage

Use `--no_hide_all` to force the compiler to use STV_DEFAULT visibility for all `extern` variables and functions if they do not use `__declspec(dll*)`. This also forces them to be preemptible at runtime by a dynamic loader.

When building a System V shared library, use `--no_hide_all` together with `--apcs /fpic`.

Use `--hide_all` to set the visibility to `STV_HIDDEN`, so that symbols cannot be dynamically linked.

Default

The default is `--hide_all`.

See also

- *--apcs=qualifer...qualifier* on page 3-7
- *__declspec(dllexport)* on page 5-19
- *__declspec(dllimport)* on page 5-21
- *--visibility_inlines_hidden* on page 3-95.

3.1.71 `-Idir[,dir,...]`

This option adds the specified directory, or comma-separated list of directories, to the list of places that are searched to find included files.

If you specify more than one directory, the directories are searched in the same order as the `-I` options specifying them.

Syntax

`-Idir[,dir,...]`

Where:

`dir[,dir,...]` is a comma-separated list of directories to be searched for included files.

At least one directory must be specified.

When specifying multiple directories, do not include spaces between commas and directory names in the list.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *-Jdir[,dir,...]* on page 3-55
- *--kandr_include* on page 3-55
- *--preinclude=filename* on page 3-78
- *--sys_include* on page 3-90
- *Compiler command-line options and search paths* on page 3-18 in *Using the Compiler*
- *Factors influencing how the compiler searches for header files* on page 3-17 in *Using the Compiler*.

3.1.72 `--ignore_missing_headers`

This option instructs the compiler to print dependency lines for header files even if the header files are missing. It only takes effect when dependency generation options (`--md` or `-M`) are specified.

Warning and error messages on missing header files are suppressed, and compilation continues.

Usage

This option is used for automatically updating makefiles. It is analogous to the GCC `-MG` command-line option.

See also

- *--depend=filename* on page 3-26
- *--depend_format=string* on page 3-27
- *--depend_system_headers, --no_depend_system_headers* on page 3-29
- *--depend_target=target* on page 3-30
- *-M* on page 3-64
- *--md* on page 3-65
- *--phony_targets* on page 3-77.

3.1.73 `--implicit_include`, `--no_implicit_include`

This option controls the implicit inclusion of source files as a method of finding definitions of template entities to be instantiated in C++.

Mode

This option is effective only if the source language is C++.

Default

The default is `--implicit_include`.

See also

- *--implicit_include_searches, --no_implicit_include_searches* on page 3-51
- *Implicit inclusion* on page 6-12.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.74 --implicit_include_searches, --no_implicit_include_searches

This option controls how the compiler searches for implicit include files for templates in C++.

When the option `--implicit_include_searches` is selected, the compiler uses the search path to look for implicit include files based on partial names of the form `filename.*`. The search path is determined by `-I`, `-J`, the `ARMCCnnINC` environment variable and the `ARMINC` environment variable. The search path also includes the default `../include` directory if `-J`, `ARMCCnnINC`, and `ARMINC` are not set.

When the option `--no_implicit_include_searches` is selected, the compiler looks for implicit include files based on the full names of files, including path names.

Mode

This option is effective only if the source language is C++.

Default

The default is `--no_implicit_include_searches`.

See also

- *-Idir[,dir,...]* on page 3-49
- *--implicit_include, --no_implicit_include* on page 3-50
- *-Jdir[,dir,...]* on page 3-55
- *Implicit inclusion* on page 6-12
- *Compiler command-line options and search paths* on page 3-18 in *Using the Compiler*
- *Toolchain environment variables* on page 2-12 in *Introducing the ARM Compiler toolchain*.

3.1.75 --implicit_key_function, --no_implicit_key_function

These options control whether an implicitly instantiated template member function can be selected as a key function. (Normally the key, or decider, function for a class is its first non-inline virtual function, in declaration order, that is not pure virtual. However, in the case of an implicitly instantiated template function, the function would have vague linkage, that is, might be multiply defined.)

Remark `#2819-D` is produced when a key function is implicit. This remark can be seen with `--remarks` or with `--diag_warning=2819`.

Default

The default is `--implicit_key_function`.

See also

- *--diag_warning=tag[,tag,...]* on page 3-33
- *--remarks* on page 3-82.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.76 `--implicit_typename, --no_implicit_typename`

This option controls the implicit determination, from context, whether a template parameter dependent name is a type or nontype in C++.

Note

The option `--implicit_typename` is provided only as a migration aid for legacy source code that does not conform to the C++ standard. Its use is not recommended.

Mode

This option is effective only if the source language is C++.

Default

The default is `--no_implicit_typename`.

Note

The `--implicit_typename` option has no effect unless you also specify `--no_parse_templates`.

See also

- *--dep_name, --no_dep_name* on page 3-25
- *--parse_templates, --no_parse_templates* on page 3-74
- *Template instantiation* on page 6-12.

3.1.77 `--import_all_vtbl`

This option causes external references to class impedimenta variables (vtables, RTTI, for example) to be marked as having dynamic linkage. It does not cause definitions of class impedimenta to have dynamic linkage.

See also

- *--export_all_vtbl, --no_export_all_vtbl* on page 3-38.

3.1.78 `--info=totals`

This option instructs the compiler to give totals of the object code and data size for each object file.

The compiler returns the same totals that `fromelf` returns when `fromelf --text -z` is used, in a similar format. The totals include embedded assembler sizes when embedded assembly exists in the source code.

Example

```
Code (inc. data)   RO Data    RW Data    ZI Data      Debug   File Name
3308       1556          0         44      10200       8402   dhry_1.o
Code (inc. data)   RO Data    RW Data    ZI Data      Debug   File Name
416          28          0          0          0       7722   dhry_2.o
```

The `(inc. data)` column gives the size of constants, string literals, and other data items used as part of the code. The `Code` column, shown in the example, *includes* this value.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--list* on page 3-59
- *--info=topic[,topic,...]* on page 2-59 in the *Linker Reference*
- *--text* on page 4-49 in *Using the fromelf Image Converter*
- *Code metrics on page 5-15* in *Using the Compiler.*

3.1.79 `--inline, --no_inline`

These options enable and disable the inlining of functions. Disabling the inlining of functions can help to improve the debug illusion.

When the option `--inline` is selected, the compiler considers inlining each function. Compiling your code with `--inline` does not guarantee that all functions are inlined. See *Compiler decisions on function inlining on page 5-30* in *Using the ARM Compiler* for more information about how the compiler decides to inline functions.

When the option `--no_inline` is selected, the compiler does not attempt to inline functions, other than functions qualified with `__forceinline`.

Default

The default is `--inline`.

See also

- *--autoinline, --no_autoinline* on page 3-14
- *--forceinline* on page 3-40
- *-Onum* on page 3-71
- *-Ospace* on page 3-72
- *-Otime* on page 3-73
- *__forceinline* on page 5-5
- *__inline* on page 5-7
- *Linker feedback during compilation* on page 3-22 in *Using the Compiler*
- *Inline functions on page 5-29* in *Using the Compiler*.

3.1.80 `--interface_enums_are_32_bit`

This option helps to provide compatibility between external code interfaces, with regard to the size of enumerated types.

Usage

It is not possible to link an object file compiled with `--enum_is_int`, with another object file that is compiled without `--enum_is_int`. The linker is unable to determine whether or not the enumerated types are used in a way that affects the external interfaces, so on detecting these build differences, it produces a warning or an error. You can avoid this by compiling with `--interface_enums_are_32_bit`. The resulting object file can then be linked with any other object file, without the linker-detected conflict that arises from different enumeration type sizes.

Note

When you use this option, you are making a promise to the compiler that all the enumerated types used in your external interfaces are 32 bits wide. For example, if you ensure that every `enum` you declare includes at least one value larger than 2 to the power of 16, the compiler is

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

forced to make the `enum` 32 bits wide, whether or not you use `--enum_is_int`. It is up to you to ensure that the promise you are making to the compiler is true. (Another method of satisfying this condition is to ensure that you have no `enums` in your external interface.)

Default

By default, the smallest data type that can hold the values of all enumerators is used.

See also

- *--enum_is_int* on page 3-36.

3.1.81 `--interleave`

This option interleaves C or C++ source code line by line as comments within an assembly listing generated using the `--asm` option or `-S` option.

Usage

The action of `--interleave` depends on the combination of options used:

Table 3-3 Compiling with the ---interleave option

Compiler option	Action
`--asm --interleave`	Writes a listing to a file of the disassembly of the compiled source, interleaving the source code with the disassembly. The link step is also performed, unless the `-c` option is used. The disassembly is written to a text file whose name defaults to the name of the input file with the filename extension `.txt`
`-S --interleave`	Writes a listing to a file of the disassembly of the compiled source, interleaving the source code with the disassembly. The disassembly is written to a text file whose name defaults to the name of the input file with the filename extension `.txt`

Restrictions

- You cannot re-assemble an assembly listing generated with `--asm --interleave` or `-S --interleave`.
- Preprocessed source files contain `#line` directives. When compiling preprocessed files using `--asm --interleave` or `-S --interleave`, the compiler searches for the original files indicated by any `#line` directives, and uses the correct lines from those files. This ensures that compiling a preprocessed file gives exactly the same output and behavior as if the original files were compiled.

  If the compiler cannot find the original files, it is unable to interleave the source. Therefore, if you have preprocessed source files with `#line` directives, but the original unpreprocessed files are not present, you must remove all the `#line` directives before you compile with `--interleave`.

See also

- *--asm* on page 3-12
- *-S* on page 3-85.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.82 -Jdir[,dir,...]

This option adds the specified directory, or comma-separated list of directories, to the list of system includes.

Warnings and remarks are suppressed, even if `--diag_error` is used.

Angle-bracketed include files are searched for first in the list of system includes, followed by any include list specified with the option `-I`.

Note

- On Windows systems, you must enclose `ARMCCnnINC` in double quotes if you specify this environment variable on the command line, because the default path defined by the variable contains spaces. For example:

```
armcc -J"%ARMCC41INC%" -c main.c
```

Syntax

`-Jdir[,dir,...]`

Where:

`dir[,dir,...]` is a comma-separated list of directories to be added to the list of system includes.

At least one directory must be specified.

When specifying multiple directories, do not include spaces between commas and directory names in the list.

See also

- *-Idir[,dir,...]* on page 3-49
- *--kandr_include*
- *--preinclude=filename* on page 3-78
- *--sys_include* on page 3-90
- *Factors influencing how the compiler searches for header files* on page 3-17 in *Using the Compiler*
- *Compiler command-line options and search paths* on page 3-18 in *Using the Compiler*
- *Toolchain environment variables* on page 2-12 in *Introducing the ARM Compiler toolchain*.

3.1.83 --kandr_include

This option ensures that Kernighan and Ritchie search rules are used for locating included files.

The current place is defined by the original source file and is not stacked.

Default

If you do not specify `--kandr_include`, Berkeley-style searching applies.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *-Idir[,dir,...]* on page 3-49
- *-Jdir[,dir,...]* on page 3-55
- *--preinclude=filename* on page 3-78
- *--sys_include* on page 3-90
- *Factors influencing how the compiler searches for header files* on page 3-17 in *Using the Compiler*
- *Compiler search rules and the current place* on page 3-19 in *Using the Compiler*.

3.1.84 `-Lopt`

This option specifies command-line options to pass to the linker when a link step is being performed after compilation. Options can be passed when creating a partially-linked object or an executable image.

Syntax

`-Lopt`

Where:

`opt` is a command-line option to pass to the linker.

Restrictions

If an unsupported Linker option is passed to it using `-L`, an error is generated by the linker.

Example

```
armcc main.c -L--map
```

See also

- *-Aopt* on page 3-6
- *--show_cmdline* on page 3-85.

3.1.85 `--library_interface=lib`

This option enables the generation of code that is compatible with the selected library type.

Syntax

`--library_interface=lib`

Where `lib` is one of:

`none` Specifies that the compiler output works with any ISO C90 library, but does not use AEABI-defined library functions unless they are required for the code to behave correctly. For example, this option suppresses the use of AEABI-defined functions that are introduced only as an optimization such as __aeabi_memcpy.

`armcc` Specifies that the compiler output works with the ARM runtime libraries in ARM Compiler 4.1 and later.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

`armcc_c90` Behaves similarly to `--library_interface=armcc`. The difference is that references in the input source code to function names that are not reserved by C90, are not modified by the compiler. Otherwise, some C99 `math.h` function names might be prefixed with `__hardfp_`, for example `__hardfp_tgamma`.

`aeabi_clib90_hardfp` Specifies that the compiler output works with any ISO C90 library compliant with the AEABI, and causes calls to the C library (including the math libraries) to call hardware floating-point library functions.

`aeabi_clib99_hardfp` Specifies that the compiler output works with any ISO C99 library compliant with the AEABI, and causes calls to the C library (including the math libraries) to call hardware floating-point library functions.

`aeabi_clib_hardfp` Specifies that the compiler output works with any ISO C library compliant with the AEABI.

Selecting the option `--library_interface=aeabi_clib_hardfp` is equivalent to specifying either `--library_interface=aeabi_clib90_hardfp` or `--library_interface=aeabi_clib99_hardfp`, depending on the choice of source language used.

The choice of source language is dependent both on the command-line options selected and on the filename suffixes used.

Causes calls to the C library (including the math libraries) to call hardware floating-point library functions.

`aeabi_glibc_hardfp` Specifies that the compiler output works with an AEABI-compliant version of the GNU C library, and causes calls to the C library (including the math libraries) to call hardware floating-point library functions.

`rvct30` Specifies that the compiler output is compatible with RVCT 3.0 runtime libraries.

`rvct30_c90` Behaves similarly to `rvct30`. In addition, specifies that the compiler output is compatible with any ISO C90 library.

`rvct31` Specifies that the compiler output is compatible with RVCT 3.1 runtime libraries.

`rvct31_c90` Behaves similarly to `rvct31`. In addition, specifies that the compiler output is compatible with any ISO C90 library.

`rvct40` Specifies that the compiler output is compatible with RVCT 4.0 runtime libraries.

`rvct40_c90` Behaves similarly to `rvct40`. In addition, specifies that the compiler output is compatible with any ISO C90 library.

Default

If you do not specify `--library_interface`, the compiler assumes `--library_interface=armcc`.

Usage

- Use the option `--library_interface=armcc` to exploit the full range of compiler and library optimizations when linking.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- Use an option of the form `--library_interface=aeabi_*` when linking with an ABI-compliant C library. Options of the form `--library_interface=aeabi_*` ensure that the compiler does not generate calls to any optimized functions provided by the ARM C library.
- It is an error to use any of the `_hardfp` library interfaces when compiling with `--fpu=softvfp`.

See also

- *Compliance with the Application Binary Interface (ABI) for the ARM architecture* on page 2-9 in *Using ARM® C and C++ Libraries and Floating-Point Support*.

3.1.86 `--library_type=lib`

This option enables the selected library to be used at link time.

Note

Use this option with the linker to override all other `--library_type` options.

Syntax

`--library_type=lib`

Where *lib* is one of:

`standardlib` Specifies that the full ARM Compiler 4.1 and later runtime libraries are selected at link time.

Use this option to exploit the full range of compiler optimizations when linking.

`microlib` Specifies that the C micro-library (microlib) is selected at link time.

Default

If you do not specify `--library_type`, the compiler assumes `--library_type=standardlib`.

See also

- *--library_type=lib* on page 2-76 in the *Linker Reference*
- *About microlib* on page 3-2 in *Using ARM® C and C++ Libraries and Floating-Point Support*
- *Building an application with microlib* on page 3-7 in *Using ARM® C and C++ Libraries and Floating-Point Support*.

3.1.87 `--link_all_input, --no_link_all_input`

This option enables and disables the suppression of errors for unrecognized input filename extensions.

When enabled, the compiler suppresses errors for unrecognized input filename extensions, and treats all unrecognized input files as object files or libraries to be passed to the linker.

Default

The default is `--no_link_all_input`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--compile_all_input, --no_compile_all_input* on page 3-20
- *Filename suffixes recognized by the compiler* on page 3-14 in *Using the Compiler*.

3.1.88 `--list`

This option instructs the compiler to generate raw listing information for a source file. The name of the raw listing file defaults to the name of the input file with the filename extension `.lst`.

If you specify multiple source files on the command line, the compiler generates listings for all of the source files, writing each to a separate listing file whose name is generated from the corresponding source file name. However, when `--multifile` is used, a concatenated listing is written to a single listing file, whose name is generated from the first source file name.

Usage

Typically, raw listing information is used to generate a formatted listing. The raw listing file contains raw source lines, information on transitions into and out of include files, and diagnostics generated by the compiler. Each line of the listing file begins with any of the following key characters that identifies the type of line:

`N` A normal line of source. The rest of the line is the text of the line of source.

`X` The expanded form of a normal line of source. The rest of the line is the text of the line. This line appears following the `N` line, and only if the line contains nontrivial modifications. Comments are considered trivial modifications, and macro expansions, line splices, and trigraphs are considered nontrivial modifications. Comments are replaced by a single space in the expanded-form line.

`S` A line of source skipped by an `#if` or similar. The rest of the line is text.

Note

The `#else`, `#elseif`, or `#endif` that ends a skip is marked with an `N`.

`L` Indicates a change in source position. That is, the line has a format similar to the `#` line-identifying directive output by the preprocessor:

`L` *`line-number`* `"`*`filename`*`"` *`key`*

where *`key`* can be:

`1` For entry into an include file.

`2` For exit from an include file.

Otherwise, *`key`* is omitted. The first line in the raw listing file is always an `L` line identifying the primary input file. `L` lines are also output for `#line` directives where *`key`* is omitted. `L` lines indicate the source position of the following source line in the raw listing file.

`R/W/E` Indicates a diagnostic, where:

`R` Indicates a remark.

`W` Indicates a warning.

`E` Indicates an error.

The line has the form:

`type` `"`*`filename`*`"` *`line-number column-number message-text`*

where *`type`* can be `R`, `W`,or `E`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Errors at the end of file indicate the last line of the primary source file and a column number of zero.

Command-line errors are errors with a filename of "`<command line>`". No line or column number is displayed as part of the error message.

Internal errors are errors with position information as usual, and message-text beginning with `(Internal fault)`.

When a diagnostic message displays a list, for example, all the contending routines when there is ambiguity on an overloaded call, the initial diagnostic line is followed by one or more lines with the same overall format. However, the code letter is the lowercase version of the code letter in the initial line. The source position in these lines is the same as that in the corresponding initial line.

Example

```
/* main.c */
#include <stdbool.h>
int main(void)
{
    return(true);
}
```

Compiling this code with the option `--list` produces the raw listing file:

```
L 1 "main.c"
N#include <stdbool.h>
L 1 "...\include\...\stdbool.h" 1
N/* stdbool.h */
N
...
N  #ifndef __cplusplus /* In C++, 'bool', 'true' and 'false' and keywords */
N    #define bool _Bool
N    #define true 1
N    #define false 0
N  #endif
...
L 2 "main.c" 2
N
Nint main(void)
N{
N   return(true);
X   return(1);
N}
```

See also

- *--asm* on page 3-12
- *-c* on page 3-17
- *--depend=filename* on page 3-26
- *--depend_format=string* on page 3-27
- *--info=totals* on page 3-52
- *--interleave* on page 3-54
- *--list_dir=directory_name* on page 3-61
- *--md* on page 3-65
- *-S* on page 3-85
- *Severity of compiler diagnostic messages* on page 6-3 in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.89 `--list_dir=directory_name`

This option enables you to specify a directory for `--list` output.

Example

```
armcc -c --list_dir=lst --list f1.c f2.c
```

Result:

```
lst/f1.lst
lst/f2.lst
```

See also

- *--asm_dir=directory_name* on page 3-13
- *--depend_dir=directory_name* on page 3-27
- *--list* on page 3-59
- *--output_dir=directory_name* on page 3-73.

3.1.90 `--list_macros`

This option lists macro definitions to `stdout` after processing a specified source file. The listed output contains macro definitions that are used on the command line, predefined by the compiler, and found in header and source files, depending on usage.

Usage

To list macros that are defined on the command line, predefined by the compiler, and found in header and source files, use `--list_macros` with a non-empty source file.

To list only macros predefined by the compiler and specified on the command line, use `--list_macros` with an empty source file.

Restrictions

Code generation is suppressed.

See also

- *Compiler predefines* on page 5-98
- *-Dname[(parm-list)][=def]* on page 3-22
- *-E* on page 3-35
- *--show_cmdline* on page 3-85
- *--via=filename* on page 3-95.

3.1.91 `--littleend`

This option instructs to the compiler to generate code for an ARM processor using little-endian memory.

With little-endian memory, the least significant byte of a word has the lowest address.

Default

The compiler assumes `--littleend` unless `--bigend` is explicitly specified.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--bigend* on page 3-14.

3.1.92 `--locale=lang_country`

This option switches the default locale for source files to the one you specify in *`lang_country`*.

Syntax

`--locale=lang_country`

Where:

`lang_country` is the new default locale.

Use this option in combination with `--multibyte_chars`.

Restrictions

The locale name might be case-sensitive, depending on the host platform.

The permitted settings of locale are determined by the host platform.

Ensure that you have installed the appropriate locale support for the host platform.

Example

To compile Japanese source files on an English-based Windows workstation, use:

```
--multibyte_chars --locale=japanese
```

and on a UNIX workstation use:

```
--multibyte_chars --locale=ja_JP
```

See also

- *--message_locale=lang_country[.codepage]* on page 3-65
- *--multibyte_chars, --no_multibyte_chars* on page 3-67.

3.1.93 `--long_long`

This option permits use of the `long long` data type in strict mode.

Example

To successfully compile the following code in strict mode, you must use `--strict --long_long`.

```
long long f(long long x, long long y)
{
    return x*y;
}
```

See also

- *--strict, --no_strict* on page 3-88.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.94 `--loose_implicit_cast`

This option makes illegal implicit casts legal, such as implicit casts of a nonzero integer to a pointer.

Example

```
int *p = 0x8000;
```

Compiling this example without the option `--loose_implicit_cast`, generates an error.

Compiling this example with the option `--loose_implicit_cast`, generates a warning message, that you can suppress.

3.1.95 `--lower_ropi, --no_lower_ropi`

This option enables or disables less restrictive C when compiling with `--apcs=/ropi`.

Default

The default is `--no_lower_ropi`.

Note

If you compile with `--lower_ropi`, then the static initialization is done at runtime by the C++ constructor mechanism for both C and C++ code. This enables these static initializations to work with ROPI code.

See also

- *--apcs=qualifer...qualifier* on page 3-7
- *--lower_rwpi, --no_lower_rwpi*
- *Code compatibility between separately compiled and assembled modules* on page 3-21 in *Using the Compiler*.

3.1.96 `--lower_rwpi, --no_lower_rwpi`

This option enables or disables less restrictive C and C++ when compiling with `--apcs=/rwpi`.

Default

The default is `--lower_rwpi`.

Note

If you compile with `--lower_rwpi`, then the static initialization is done at runtime by the C++ constructor mechanism, even for C. This enables these static initializations to work with RWPI code.

See also

- *--apcs=qualifer...qualifier* on page 3-7
- *--lower_ropi, --no_lower_ropi*
- *Code compatibility between separately compiled and assembled modules* on page 3-21 in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.97 `--ltcg`

This option instructs the compiler to create objects in an intermediate format so that *Link-Time Code Generation* (LTCG) optimizations can be performed. The optimizations applied include cross-module inlining to improve performance, and sharing of base addresses to reduce code size.

Note

- This option might significantly increase link time and memory requirements. For large applications it is recommended that you do the code generation in partial link steps with a subset of the objects.
- The LTCG feature is deprecated. As an alternative ARM recommends you use the `--multifile` option.

Example

The following example shows how to use the `--ltcg` option.

```
armcc -c --ltcg file1.c
armcc -c --ltcg file2.c
armlink --ltcg file1.o file2.o -o prog.axf
```

See also

- *--multifile, --no_multifile* on page 3-67
- *-Onum* on page 3-71
- *Automatic function inlining and multifile compilation on page 5-34* in Using the Compiler
- *Inline functions in C99 mode on page 5-38* in Using the Compiler
- *--ltcg* on page 2-81 in the *Linker Reference*
- *About link-time code generation* on page 5-10 in the *Linker Reference*.

3.1.98 `-M`

This option instructs the compiler to produce a list of makefile dependency lines suitable for use by a make utility.

The compiler executes only the preprocessor step of the compilation. By default, output is on the standard output stream.

If you specify multiple source files, a single dependency file is created.

If you specify the `-o` *`filename`* option, the dependency lines generated on standard output make reference to `filename.o`, and not to `source.o`. However, no object file is produced with the combination of `-M -o` *`filename`*.

Use the `--md` option to generate dependency lines and object files for each source file.

Example

You can redirect output to a file by using standard UNIX and MS-DOS notation, for example:

```
armcc -M source.c > Makefile
```

See also

- *-C* on page 3-17

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *--depend=filename* on page 3-26
- *--depend_system_headers, --no_depend_system_headers* on page 3-29
- *-E* on page 3-35
- *--md*
- *--depend_single_line, --no_depend_single_line* on page 3-28
- *-o filename* on page 3-69.

3.1.99 `--md`

This option instructs the compiler to compile the source and write makefile dependency lines to a file.

The output file is suitable for use by a make utility.

The compiler names the file `*filename*.d`, where `*filename*` is the name of the source file. If you specify multiple source files, a dependency file is created for each source file.

If you want to produce makefile dependencies and preprocessor source file output in a single step, you can do so using the combination `--md -E` (or `--md -P` to suppress line number generation).

See also

- *--depend=filename* on page 3-26
- *--depend_format=string* on page 3-27
- *--depend_system_headers, --no_depend_system_headers* on page 3-29
- *-E* on page 3-35
- *-M* on page 3-64
- *--depend_single_line, --no_depend_single_line* on page 3-28
- *-o filename* on page 3-69.

3.1.100 `--message_locale=*lang_country*[.*codepage*]`

This option switches the default language for the display of error and warning messages to the one you specify in `*lang_country*` or `*lang_country.codepage*`.

Syntax

`--message_locale=*lang_country*[.*codepage*]`

Where:

`*lang_country*[.*codepage*]`

is the new default language for the display of error and warning messages.

The permitted languages are independent of the host platform.

The following settings are supported:

- `en_US`
- `zh_CN`
- `ko_KR`
- `ja_JP`.

Default

If you do not specify `--message_locale`, the compiler assumes `--message_locale=en_US`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Restrictions

Ensure that you have installed the appropriate locale support for the host platform.

The locale name might be case-sensitive, depending on the host platform.

The ability to specify a codepage, and its meaning, depends on the host platform.

Errors

If you specify a setting that is not supported, the compiler generates an error message.

Example

To display messages in Japanese, use:

```
--message_locale=ja_JP
```

See also

- *--locale=lang_country* on page 3-62
- *--multibyte_chars, --no_multibyte_chars* on page 3-67.

3.1.101 `--min_array_alignment=opt`

This option enables you to specify the minimum alignment of arrays.

Syntax

`--min_array_alignment=opt`

Where:

opt specifies the minimum alignment of arrays. The value of `opt` is:

1	byte alignment, or unaligned
2	two-byte, halfword alignment
4	four-byte, word alignment
8	eight-byte, doubleword alignment.

Usage

Use of this option is not recommended, unless required in certain specialized cases. For example, porting code to systems that have different data alignment requirements. Use of this option can result in increased code size at the higher `opt` values, and reduced performance at the lower `opt` values. If you only want to affect the alignment of specific arrays (rather than all arrays), use the **`__align`** keyword instead.

Default

If you do not use this option, arrays are unaligned (byte aligned).

Example

Compiling the following code with `--min_array_alignment=8` gives the alignment described in the comments:

```
char arr_c1[1];      // alignment == 8
char c1;             // alignment == 1
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *__align* on page 5-2
- *__ALIGNOF__* on page 5-3.

3.1.102 `--mm`

This option has the same effect as `-M --no_depend_system_headers`.

See also

- *--depend_system_headers, --no_depend_system_headers* on page 3-29
- *-M* on page 3-64.

3.1.103 `--multibyte_chars`, `--no_multibyte_chars`

This option enables or disables processing for multibyte character sequences in comments, string literals, and character constants.

Default

The default is `--no_multibyte_chars`.

Usage

Multibyte encodings are used for character sets such as the Japanese *Shift-Japanese Industrial Standard* (Shift-JIS).

See also

- *--locale=lang_country* on page 3-62
- *--message_locale=lang_country[.codepage]* on page 3-65.

3.1.104 `--multifile`, `--no_multifile`

This option enables or disables multifile compilation.

When `--multifile` is selected, the compiler performs optimizations across all files specified on the command line, instead of on each individual file. The specified files are compiled into one single object file.

The combined object file is named after the first source file you specify on the command line. To specify a different name for the combined object file, use the `-o filename` option.

An empty object file is created for each subsequent source file specified on the command line to meet the requirements of standard make systems.

Note

Compiling with `--multifile` has no effect if only a single source file is specified on the command line.

Default

The default is `--no_multifile`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Usage

When `--multifile` is selected, the compiler might be able to perform additional optimizations by compiling across several source files.

There is no limit to the number of source files that can be specified on the command line, but ten files is a practical limit, because `--multifile` requires large amounts of memory while compiling. For the best optimization results, choose small groups of functionally related source files.

Example

```
armcc -c --multifile test1.c ... testn.c -o test.o
```

The resulting object file is named `test.o`, instead of `test1.c`, and empty object files `test2.o` to `testn.o` are created for each source file `test1.c ... testn.c` specified on the command line.

See also

- *-c* on page 3-17
- *--default_extension=ext* on page 3-25
- *--ltcg* on page 3-64
- *-o filename* on page 3-69
- *-Onum* on page 3-71
- *--whole_program* on page 3-98
- *Predefined macros* on page 5-98.

3.1.105 `--multiply_latency=cycles`

This option tells the compiler the number of cycles used by the hardware multiplier.

Syntax

```
--multiply_latency=cycles
```

Where *`cycles`* is the number of cycles used.

Usage

Use this option to tell the compiler how many cycles the MUL instruction takes to use the multiplier block and related parts of the chip. Until finished, these parts of the chip cannot be used for another instruction and the result of the MUL is not available for any later instructions to use.

It is possible that a processor might have two or more multiplier options that are set for a given hardware implementation. For example, one implementation might be configured to take one cycle to execute. The other implementation might take 33 cycles to execute. This option is used to convey the correct number of cycles for a given processor.

Default

The default number of cycles used by the hardware multiplier is processor-specific. See the Technical Reference Manual for the processor architecture you are compiling for.

Example

```
--multiply_latency=33
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *Cortex™-M1 Technical Reference Manual*.

3.1.106 `--narrow_volatile_bitfields`

The AEABI specifies that volatile bitfields are accessed as the size of their container type. However, some versions of GCC instead use the smallest access size that contains the entire bitfield. `--narrow_volatile_bitfields` emulates this non-AEABI compliant behavior.

See also

- *Application Binary Interface* (ABI) for the ARM Architecture, http://infocenter.arm.com/help/index.jsp?topic=/com.arm.doc.subset.swdev.abi/index.html

3.1.107 `--nonstd_qualifier_deduction`, `--no_nonstd_qualifier_deduction`

This option controls whether or not nonstandard template argument deduction is to be performed in the qualifier portion of a qualified name in C++.

With this feature enabled, a template argument for the template parameter `T` can be deduced in contexts like `A<T>::B` or `T::B`. The standard deduction mechanism treats these as nondeduced contexts that use the values of template parameters that were either explicitly specified or deduced elsewhere.

Note

The option `--nonstd_qualifier_deduction` is provided only as a migration aid for legacy source code that does not conform to the C++ standard. Its use is not recommended.

Mode

This option is effective only if the source language is C++.

Default

The default is `--no_nonstd_qualifier_deduction`.

3.1.108 `-o` *`filename`*

This option specifies the name of the output file. The full name of the output file produced depends on the combination of options used, as described in Table 3-4 on page 3-70 and Table 3-5 on page 3-70.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Syntax

If you specify a `-o` option, the compiler names the output file according to the conventions of Table 3-4.

Table 3-4 Compiling with the -o option

Compiler option	Action	Usage notes
`-o-`	writes output to the standard output stream	*filename* is `-`.`-S` is assumed unless `-E` is specified.
`-o` *filename*	produces an executable image with name *filename*	
`-c -o` *filename*	produces an object file with name *filename*	
`-S -o` *filename*	produces an assembly language file with name *filename*	
`-E -o` *filename*	produces a file containing preprocessor output with name *filename*	

Note

This option overrides the `--default_extension` option.

Default

If you do not specify a `-o` option, the compiler names the output file according to the conventions of Table 3-5.

Table 3-5 Compiling without the -o option

Compiler option	Action	Usage notes
`-c`	produces an object file whose name defaults to the name of the input file with the filename extension `.o`	
`-S`	produces an output file whose name defaults to the name of the input file with the filename extension `.s`	
`-E`	writes output from the preprocessor to the standard output stream	
(No option)	produces an executable image with the default name of `__image.axf`	none of `-o`, `-c`, `-E` or `-S` is specified on the command line

See also

- *--asm* on page 3-12
- *-c* on page 3-17
- *--default_extension=ext* on page 3-25
- *--depend=filename* on page 3-26
- *--depend_format=string* on page 3-27

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *-E* on page 3-35
- *--interleave* on page 3-54
- *--list* on page 3-59
- *--md* on page 3-65
- *-S* on page 3-85.

3.1.109 `-Onum`

This option specifies the level of optimization to be used when compiling source files.

Syntax

`-Onum`

Where *num* is one of the following:

`0` Minimum optimization. Turns off most optimizations. It gives the best possible debug view and the lowest level of optimization.

`1` Restricted optimization. Removes unused inline functions and unused static functions. Turns off optimizations that seriously degrade the debug view. If used with `--debug`, this option gives a satisfactory debug view with good code density.

`2` High optimization. If used with `--debug`, the debug view might be less satisfactory because the mapping of object code to source code is not always clear.

This is the default optimization level.

`3` Maximum optimization. `-O3` performs the same optimizations as `-O2` however the balance between space and time optimizations in the generated code is more heavily weighted towards space or time compared with `-O2`. That is:

- `-O3 -Otime` aims to produce faster code than `-O2 -Otime`, at the risk of increasing your image size
- `-O3 -Ospace` aims to produce smaller code than `-O2 -Ospace`, but performance might be degraded.

In addition, `-O3` performs extra optimizations that are more aggressive, such as:

- High-level scalar optimizations, including loop unrolling, for `-O3 -Otime`. This can give significant performance benefits at a small code size cost, but at the risk of a longer build time.
- More aggressive inlining and automatic inlining for `-O3 -Otime`.

Note

The performance of floating-point code can be influenced by selecting an appropriate numerical model using the `--fpmode` option.

Note

Do not rely on the implementation details of these optimizations, because they might change in future releases.

Default

If you do not specify `-Onum`, the compiler assumes `-O2`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--autoinline, --no_autoinline* on page 3-14
- *--debug, --no_debug* on page 3-24
- *--forceinline* on page 3-40
- *--fpmode=model* on page 3-42
- *--inline, --no_inline* on page 3-53
- *--ltcg* on page 3-64
- *--multifile, --no_multifile* on page 3-67
- *-Ospace*
- *-Otime* on page 3-73
- *The compiler as an optimizing compiler on page 5-5* in *Using the Compiler*.

3.1.110 `--old_specializations, --no_old_specializations`

This option controls the acceptance of old-style template specializations in C++.

Old-style template specializations do not use the `template<>` syntax.

Note

The option `--old_specializations` is provided only as a migration aid for legacy source code that does not conform to the C++ standard. Its use is not recommended.

Mode

This option is effective only if the source language is C++.

Default

The default is `--no_old_specializations`.

3.1.111 `--old_style_preprocessing`

This option performs preprocessing in the style of legacy compilers that do not follow the ISO C Standard.

See also

- *-E* on page 3-35.

3.1.112 `-Ospace`

This option instructs the compiler to perform optimizations to reduce image size at the expense of a possible increase in execution time.

Use this option if code size is more critical than performance. For example, when the `-Ospace` option is selected, large structure copies are done by out-of-line function calls instead of inline code.

If required, you can compile the time-critical parts of your code with `-Otime`, and the rest with `-Ospace`.

Default

If you do not specify either `-Ospace` or `-Otime`, the compiler assumes `-Ospace`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *-Otime*
- *-Onum* on page 3-71
- *#pragma Onum* on page 5-54
- *#pragma Ospace* on page 5-55
- *#pragma Otime* on page 5-55.

3.1.113 `-Otime`

This option instructs the compiler to perform optimizations to reduce execution time at the expense of a possible increase in image size.

Use this option if execution time is more critical than code size. If required, you can compile the time-critical parts of your code with `-Otime`, and the rest with `-Ospace`.

Default

If you do not specify `-Otime`, the compiler assumes `-Ospace`.

Example

When the `-Otime` option is selected, the compiler compiles:

```
while (expression) body;
```

as:

```
if (expression)
{
    do body;
    while (expression);
}
```

See also

- *--multifile, --no_multifile* on page 3-67
- *-Onum* on page 3-71
- *-Ospace* on page 3-72
- *#pragma Onum* on page 5-54
- *#pragma Ospace* on page 5-55
- *#pragma Otime* on page 5-55.

3.1.114 `--output_dir=`*`directory_name`*

This option enables you to specify an output directory for object files and depending on the other options you use, certain other types of compiler output.

The directory for assembler output can be specified using `--asm_dir`. The directory for dependency output can be specified using `--depend_dir`. The directory for `--list` output can be specified using `--list_dir`. If these options are not used, the corresponding output is placed in the directory specified by `--output_dir`, or if `--output_dir` is not specified, in the default location (for example, the current directory).

The executable is placed in the default location.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example

```
armcc -c --output_dir=obj f1.c f2.c
```

Result:

```
obj/f1.o
obj/f2.o
```

See also

- *--asm_dir=directory_name* on page 3-13
- *--depend_dir=directory_name* on page 3-27
- *--list_dir=directory_name* on page 3-61.

3.1.115 `-P`

This option preprocesses source code without compiling, but does not generate line markers in the preprocessed output.

Usage

This option can be of use when the preprocessed output is destined to be parsed by a separate script or utility.

See also

- *-E* on page 3-35.

3.1.116 `--parse_templates`, `--no_parse_templates`

This option enables or disables the parsing of nonclass templates in their generic form in C++, that is, when the template is defined and before it is instantiated.

Note

The option `--no_parse_templates` is provided only as a migration aid for legacy source code that does not conform to the C++ standard. Its use is not recommended.

Mode

This option is effective only if the source language is C++.

Default

The default is `--parse_templates`.

Note

`--no_parse_templates` cannot be used with `--dep_name`, because parsing is done by default if dependent name processing is enabled. Combining these options generates an error.

See also

- *--dep_name, --no_dep_name* on page 3-25
- *Template instantiation* on page 6-12.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.117 `--pch`

This option instructs the compiler to use a PCH file if it exists, and to create a PCH file otherwise.

When the option `--pch` is specified, the compiler searches for a PCH file with the name *filename*.pch, where *filename*.* is the name of the primary source file. The compiler uses the PCH file *filename*.pch if it exists, and creates a PCH file named *filename*.pch in the same directory as the primary source file otherwise.

Restrictions

This option has no effect if you include either the option `--use_pch=`*filename* or the option `--create_pch=`*filename* on the same command line.

See also

- *--create_pch=filename* on page 3-22
- *--pch_dir=dir*
- *--pch_messages, --no_pch_messages* on page 3-76
- *--pch_verbose, --no_pch_verbose* on page 3-76
- *--use_pch=filename* on page 3-93
- *#pragma hdrstop* on page 5-52
- *#pragma no_pch* on page 5-54
- *PreCompiled Header (PCH) files* on page 4-29 in *Using the Compiler*.

3.1.118 `--pch_dir=`*dir*

This option enables you to specify the directory where PCH files are stored. The directory is accessed whenever PCH files are created or used.

You can use this option with automatic or manual PCH mode.

Syntax

`--pch_dir=`*dir*

Where:

dir is the name of the directory where PCH files are stored.

If *dir* is unspecified, the compiler faults use of `--pch_dir`.

Errors

If the specified directory *dir* does not exist, the compiler generates an error.

See also

- *--create_pch=filename* on page 3-22
- *--pch*
- *--pch_messages, --no_pch_messages* on page 3-76
- *--pch_verbose, --no_pch_verbose* on page 3-76
- *--use_pch=filename* on page 3-93
- *#pragma hdrstop* on page 5-52
- *#pragma no_pch* on page 5-54

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *PreCompiled Header (PCH) files* on page 4-29 in *Using the Compiler*.

3.1.119 `--pch_messages`, `--no_pch_messages`

This option enables or disables the display of messages indicating that a PCH file is used in the current compilation.

Default

The default is `--pch_messages`.

See also

- *--create_pch=filename* on page 3-22
- *--pch* on page 3-75
- *--pch_dir=dir* on page 3-75
- *--pch_verbose, --no_pch_verbose*
- *--use_pch=filename* on page 3-93
- *#pragma hdrstop* on page 5-52
- *#pragma no_pch* on page 5-54
- *PreCompiled Header (PCH) files* on page 4-29 in *Using the Compiler*.

3.1.120 `--pch_verbose`, `--no_pch_verbose`

This option enables or disables the display of messages giving reasons why a file cannot be precompiled.

In automatic PCH mode, this option ensures that for each PCH file that cannot be used for the current compilation, a message is displayed giving the reason why the file cannot be used.

Default

The default is `--no_pch_verbose`.

See also

- *--create_pch=filename* on page 3-22
- *--pch* on page 3-75
- *--pch_dir=dir* on page 3-75
- *--pch_messages, --no_pch_messages*
- *--use_pch=filename* on page 3-93
- *#pragma hdrstop* on page 5-52
- *#pragma no_pch* on page 5-54
- *PreCompiled Header (PCH) files* on page 4-29 in *Using the Compiler*.

3.1.121 `--pending_instantiations=n`

This option specifies the maximum number of concurrent instantiations of a template in C++.

Syntax

```
--pending_instantiations=n
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Where:

`n` is the maximum number of concurrent instantiations permitted.

If `n` is zero, there is no limit.

Mode

This option is effective only if the source language is C++.

Default

If you do not specify a `--pending_instantations` option, then the compiler assumes `--pending_instantiations=64`.

Usage

Use this option to detect runaway recursive instantiations.

3.1.122 --phony_targets

This option instructs the compiler to emit dummy makefile rules. These rules work around `make` errors that are generated if you remove header files without a corresponding update to the makefile.

This option is analogous to the GCC command-line option, `-MP`.

Example

Example output:

```
source.o: source.c
source.o: header.h
header.h:
```

See also

- *--depend=filename* on page 3-26
- *--depend_format=string* on page 3-27
- *--depend_system_headers, --no_depend_system_headers* on page 3-29
- *--depend_target=target* on page 3-30
- *--ignore_missing_headers* on page 3-50
- *-M* on page 3-64
- *--md* on page 3-65

3.1.123 --pointer_alignment=*num*

This option specifies the unaligned pointer support required for an application.

Syntax

`--pointer_alignment=num`

Where `num` is one of:

`1` Treats accesses through pointers as having an alignment of one, that is, byte-aligned or unaligned.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

`2` Treats accesses through pointers as having an alignment of at most two, that is, at most halfword aligned.

`4` Treats accesses through pointers as having an alignment of at most four, that is, at most word aligned.

`8` Accesses through pointers have normal alignment, that is, at most doubleword aligned.

If *num* is unspecified, the compiler faults use of `--pointer_alignment`.

Usage

This option can help you port source code that has been written for architectures without alignment requirements. You can achieve finer control of access to unaligned data, with less impact on the quality of generated code, using the `__packed` qualifier.

Restrictions

De-aligning pointers might increase the code size, even on CPUs with unaligned access support. This is because only a subset of the load and store instructions benefit from unaligned access support. The compiler is unable to use multiple-word transfers or coprocessor-memory transfers, including hardware floating-point loads and stores, directly on unaligned memory objects.

Note

- Code size might increase significantly when compiling for CPUs without hardware support for unaligned access, for example, pre-v6 architectures.
- This option does not affect the placement of objects in memory, nor the layout and padding of structures.

See also

- *__packed* on page 5-9
- *#pragma pack(n)* on page 5-56
- *Compiler storage of data objects by natural byte alignment on page 5-43* in *Using the Compiler*.

3.1.124 `--preinclude=filename`

This option instructs the compiler to include the source code of the specified file at the beginning of the compilation.

Syntax

`--preinclude=filename`

Where:

`filename` is the name of the file whose source code is to be included.

If *filename* is unspecified, the compiler faults use of `--preinclude`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Usage

This option can be used to establish standard macro definitions. The *filename* is searched for in the directories on the include search list.

It is possible to repeatedly specify this option on the command line. This results in pre-including the files in the order specified.

Example

```
armcc --preinclude file1.h --preinclude file2.h -c source.c
```

See also

- *-Idir[,dir,...]* on page 3-49
- *-Jdir[,dir,...]* on page 3-55
- *--kandr_include* on page 3-55
- *--sys_include* on page 3-90
- *Factors influencing how the compiler searches for header files* on page 3-17 in *Using the Compiler*.

3.1.125 `--preprocess_assembly`

This option relaxes certain rules when producing preprocessed compiler output, to provide greater flexibility when preprocessing assembly language source code.

Usage

Use this option to relax certain preprocessor rules when generating preprocessed output from assembly language source files. Specifically, the following special cases are permitted that would normally produce a compiler error:

- Lines beginning with a '#' character followed by a space and a number, that would normally indicate a GNU non-standard line marker, are ignored and copied verbatim into the preprocessed output.
- Unrecognized preprocessing directives are ignored and copied verbatim into the preprocessed output.
- Where the token-paste '#' operator is used in a function-like macro, if it is used with a name that is not a macro parameter, the name is copied verbatim into the preprocessed output together with the preceding '#' character.

  For example if the source file contains:

  ```
  # define mymacro(arg) foo #bar arg
  mymacro(x)
  ```

  using the `--preprocess_assembly` option produces a preprocessed output that contains:

  ```
  foo #bar x
  ```

Restrictions

This option is only valid when producing preprocessed output without continuing compilation, for example when using the -E, -P or -C command line options. It is ignored in other cases.

See also

- *-C* on page 3-17

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *-E* on page 3-35
- *-P* on page 3-74.

3.1.126 `--preprocessed`

This option forces the preprocessor to handle files with `.i` filename extensions as if macros have already been substituted.

Usage

This option gives you the opportunity to use a different preprocessor. Generate your preprocessed code and then give the preprocessed code to the compiler in the form of a `filename.i` file, using `--preprocessed` to inform the compiler that the file has already been preprocessed.

Restrictions

This option only applies to macros. Trigraphs, line concatenation, comments and all other preprocessor items are preprocessed by the preprocessor in the normal way.

If you use `--compile_all_input`, the `.i` file is treated as a `.c` file. The preprocessor behaves as if no prior preprocessing has occurred.

Example

```
armcc --preprocessed foo.i -c -o foo.o
```

See also

- *--compile_all_input, --no_compile_all_input* on page 3-20
- *-E* on page 3-35.

3.1.127 `--reassociate_saturation`, `--no_reassociate_saturation`

These options enable and disable more aggressive optimization in loops that use saturating addition, by either permitting or prohibiting re-association of saturation arithmetic.

Usage

Although potentially useful when vectorizing code, these options are not necessarily restricted to vectorization. For example, `--reassociate_saturation` could take effect when compiling with `-O3 -Otime`, even when automatic vectorization is not enabled.

Restriction

Saturating addition is not associative, so enabling re-association could affect the result with a reduction in accuracy.

Default

The default is `--no_reassociate_saturation`.

Example

The following code does not vectorize unless `--reassociate_saturation` is specified.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

```
#include <dspfns.h>
int f(short *a, short *b)
{
    int i;
    int r = 0;
    for (i = 0; i < 100; i++)
         r=L_mac(r,a[i],b[i]);
    return r;
}
```

3.1.128 `--reduce_paths`, `--no_reduce_paths`

This option enables or disables the elimination of redundant path name information in file paths.

When elimination of redundant path name information is enabled, the compiler removes sequences of the form `xyz\..` from directory paths passed to the operating system. This includes system paths constructed by the compiler itself, for example, for `#include` searching.

Note

The removal of sequences of the form `xyz\..` might not be valid if `xyz` is a link.

Mode

This option is effective on Windows systems only.

Usage

Windows systems impose a 260 character limit on file paths. Where path names exist whose absolute names expand to longer than 260 characters, you can use the `--reduce_paths` option to reduce absolute path name length by matching up directories with corresponding instances of `..` and eliminating the directory/`..` sequences in pairs.

Note

It is recommended that you avoid using long and deeply nested file paths, in preference to minimizing path lengths using the `--reduce_paths` option.

Default

The default is `--no_reduce_paths`.

Example

Compiling the file

`..\..\..\xyzzy\xyzzy\objects\file.c`

from the directory

`\foo\bar\baz\gazonk\quux\bop`

results in an actual path of

`\foo\bar\baz\gazonk\quux\bop\..\..\..\xyzzy\xyzzy\objects\file.o`

Compiling the same file from the same directory using the option `--reduce_paths` results in an actual path of

`\foo\bar\baz\xyzzy\xyzzy\objects\file.c`

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.129 `--relaxed_ref_def`, `--no_relaxed_ref_def`

This option permits multiple object files to use tentative definitions of global variables. Some traditional programs are written using this declaration style.

Usage

This option is primarily provided for compatibility with GNU C. It is not recommended for new application code.

Default

The default is strict references and definitions. (Each global variable can only be declared in one object file.)

Restrictions

This option is not available in C++.

See also

- *Rationale for International Standard - Programming Languages - C*.

3.1.130 `--remarks`

This option instructs the compiler to issue remark messages, such as warning of padding in structures.

Default

By default, the compiler does not issue remarks.

See also

- *--brief_diagnostics, --no_brief_diagnostics* on page 3-15
- *--diag_error=tag[,tag,...]* on page 3-30
- *--diag_remark=tag[,tag,...]* on page 3-31
- *--diag_style={arm|ide|gnu}* on page 3-31
- *--diag_suppress=tag[,tag,...]* on page 3-32
- *--diag_warning=tag[,tag,...]* on page 3-33
- *--errors=filename* on page 3-36
- *-W* on page 3-96
- *--wrap_diagnostics, --no_wrap_diagnostics* on page 3-98.

3.1.131 `--remove_unneeded_entities, --no_remove_unneeded_entities`

These options control whether debug information is generated for all source symbols, or only for those source symbols actually used.

Usage

Use `--remove_unneeded_entities` to reduce the amount of debug information in an ELF object. Faster linkage times can also be achieved.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Caution

Although `--remove_unneeded_entities` can help to reduce the amount of debug information generated per file, it has the disadvantage of reducing the number of debug sections that are common to many files. This reduces the number of common debug sections that the linker is able to remove at final link time, and can result in a final debug image that is larger than necessary. For this reason, use `--remove_unneeded_entities` only when necessary.

Restrictions

The effects of these options are restricted to debug information.

Default

The default is `--no_remove_unneeded_entities`.

See also

- *The DWARF Debugging Standard*, `http://dwarfstd.org/`

3.1.132 `--restrict`, `--no_restrict`

This option enables or disables the use of the C99 keyword restrict.

Note

The alternative keywords `__restrict` and `__restrict__` are supported as synonyms for **`restrict`**. These alternative keywords are always available, regardless of the use of the `--restrict` option.

Default

When compiling ISO C99 source code, use of the C99 keyword **`restrict`** is enabled by default.

When compiling ISO C90 or ISO C++ source code, use of the C99 keyword **`restrict`** is disabled by default.

See also

- *restrict* on page 4-7.

3.1.133 `--retain=`*`option`*

This option enables you to restrict the optimizations performed by the compiler.

Syntax

`--retain=`*`option`*

Where *`option`* is one of the following:

`fns` prevents the removal of unused functions

`inlinefns` prevents the removal of unused inline functions

`noninlinefns` prevents the removal of unused non-inline functions

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

`paths` prevents path-removing optimizations, such as `a||b` transformed to `a|b`. This supports *Modified Condition Decision Coverage* (MCDC) testing.

`calls` prevents calls being removed, for example by inlining or tailcalling.

`calls:distinct`
prevents calls being merged, for example by cross-jumping (that is, common tail path merging).

`libcalls` prevents calls to library functions being removed, for example by inline expansion.

`data` prevents data being removed.

`rodata` prevents read-only data being removed.

`rwdata` prevents read-write data being removed.

`data:order` prevents data being reordered.

If *`option`* is unspecified, the compiler faults use of `--retain`.

Usage

This option might be useful when performing validation, debugging, and coverage testing. In most other cases, it is not required.

Using this option can have a negative effect on code size and performance.

See also

- *__attribute__((nomerge)) function attribute* on page 5-30
- *__attribute__((notailcall)) function attribute* on page 5-31.

3.1.134 `--rtti, --no_rtti`

This option controls support for the RTTI features `dynamic_cast` and `typeid` in C++.

Note

You are permitted to use **`dynamic_cast`** without `--rtti` in cases where RTTI is not required, such as dynamic cast to an unambiguous base, and dynamic cast to `(void *)`. If you try to use **`dynamic_cast`** without `--rtti` in cases where RTTI *is* required, the compiler generates an error.

Mode

This option is effective only if the source language is C++.

Default

The default is `--rtti`.

See also

- *--rtti_data, --no_rtti_data*.

3.1.135 `--rtti_data, --no_rtti_data`

These options enable and disable the generation of C++ RTTI data.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Note

The option `--no_rtti` only disables source-level RTTI features such as `dynamic_cast`, whereas `--no_rtti_data` disables both source-level features and the generation of RTTI data.

Mode

This option is effective only if the source language is C++.

Default

The default is `--rtti_data`.

See also

- *--rtti, --no_rtti* on page 3-84

3.1.136 `-S`

This option instructs the compiler to output the disassembly of the machine code generated by the compiler to a file.

Unlike the `--asm` option, object modules are not generated. The name of the assembly output file defaults to *`filename`*`.s` in the current directory, where *`filename`* is the name of the source file stripped of any leading directory names. The default filename can be overridden with the `-o` option.

You can use `armasm` to assemble the output file and produce object code. The compiler adds `ASSERT` directives for command-line options such as AAPCS variants and byte order to ensure that compatible compiler and assembler options are used when re-assembling the output. You must specify the same AAPCS settings to both the assembler and the compiler.

See also

- *--apcs=qualifer...qualifier* on page 3-7
- *--asm* on page 3-12
- *-c* on page 3-17
- *--info=totals* on page 3-52
- *--interleave* on page 3-54
- *--list* on page 3-59
- *-o filename* on page 3-69

- *Assembler Guide.*

3.1.137 `--show_cmdline`

This option shows how the compiler processes the command line. It can be useful when checking:

- the command line a build system is using
- how the compiler is interpreting the supplied command line, for example, the ordering of command-line options.

The commands are shown normalized, and the contents of any via files are expanded.

The output is sent to the standard output stream (`stdout`).

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *-Aopt* on page 3-6
- *-Lopt* on page 3-56
- *--via=filename* on page 3-95.

3.1.138 `--signed_bitfields`, `--unsigned_bitfields`

This option makes bitfields of type **`int`** signed or unsigned.

The C Standard specifies that if the type specifier used in declaring a bitfield is either **`int`**, or a **`typedef`** name defined as **`int`**, then whether the bitfield is signed or unsigned is dependent on the implementation.

Default

The default is `--unsigned_bitfields`.

Note

The AAPCS requirement for bitfields to default to unsigned on ARM, is relaxed in version 2.03 of the standard.

Example

```
typedef int integer;
struct
{
    integer x : 1;
} bf;
```

Compiling this code with `--signed_bitfields` causes to be treated as a signed bitfield.

See also

- *Procedure Call Standard for the ARM® Architecture*, http://infocenter/help/index.jsp?topic=/com.arm.doc.ihi0042-

3.1.139 `--signed_chars`, `--unsigned_chars`

This option makes the **`char`** type signed or unsigned.

When **`char`** is signed, the macro `__FEATURE_SIGNED_CHAR` is also defined by the compiler.

Note

Care must be taken when mixing translation units that have been compiled with and without this option, and that share interfaces or data structures.

The ARM ABI defines **`char`** as an unsigned byte, and this is the interpretation used by the C++ libraries.

Default

The default is `--unsigned_chars`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *Predefined macros* on page 5-98.

3.1.140 `--split_ldm`

This option instructs the compiler to split `LDM` and `STM` instructions into two or more `LDM` or `STM` instructions.

When `--split_ldm` is selected, the maximum number of register transfers for an `LDM` or `STM` instruction is limited to:

- five, for all `STM`s
- five, for `LDM`s that do not load the PC
- four, for `LDM`s that load the PC.

Where register transfers beyond these limits are required, multiple `LDM` or `STM` instructions are used.

Usage

The `--split_ldm` option can be used to reduce interrupt latency on ARM systems that:

- do not have a cache or a write buffer, for example, a cacheless ARM7TDMI
- use zero-wait-state, 32-bit memory.

Note

Using `--split_ldm` increases code size and decreases performance slightly.

Restrictions

- Inline assembler `LDM` and `STM` instructions are split by default when `--split_ldm` is used. However, the compiler might subsequently recombine the separate instructions into an `LDM` or `STM`.
- Only `LDM` and `STM` instructions are split when `--split_ldm` is used.
- Some target hardware does not benefit from code built with `--split_ldm`. For example:
  - It has no significant benefit for cached systems, or for processors with a write buffer.
  - It has no benefit for systems with non zero-wait-state memory, or for systems with slow peripheral devices. Interrupt latency in such systems is determined by the number of cycles required for the slowest memory or peripheral access. Typically, this is much greater than the latency introduced by multiple register transfers.

See also

- *Inline assembler and instruction expansion in C and C++ code* on page 7-21 in *Using the Compiler*.

3.1.141 `--split_sections`

This option instructs the compiler to generate one ELF section for each function in the source file.

Output sections are named with the same name as the function that generates the section, but with an `i.` prefix.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Note

If you want to place specific data items or structures in separate sections, mark them individually with `__attribute__((section(...)))`.

If you want to remove unused functions, it is recommended that you use the linker feedback optimization in preference to this option. This is because linker feedback produces smaller code by avoiding the overhead of splitting all sections.

Restrictions

This option reduces the potential for sharing addresses, data, and string literals between functions. Consequently, it might increase code size slightly for some functions.

Example

```
int f(int x)
{
    return x+1;
}
```

Compiling this code with `--split_sections` produces:

```
        AREA ||i.f||, CODE, READONLY, ALIGN=2
f PROC
        ADD      r0,r0,#1
        BX       lr
        ENDP
```

See also

- *--data_reorder, --no_data_reorder* on page 3-23
- *--feedback=filename* on page 3-39
- *--multifile, --no_multifile* on page 3-67
- *__attribute__((section("name"))) function attribute* on page 5-32
- *#pragma arm section [section_type_list]* on page 5-48
- *Linker feedback during compilation* on page 3-22 in *Using the Compiler*.

3.1.142 `--strict, --no_strict`

This option enforces or relaxes strict C or strict C++, depending on the choice of source language used.

When `--strict` is selected:

- features that conflict with ISO C or ISO C++ are disabled
- error messages are returned when nonstandard features are used.

Default

The default is `--no_strict`.

Usage

`--strict` enforces compliance with:

ISO C90
- ISO/IEC 9899:1990, the 1990 International Standard for C.
- ISO/IEC 9899 AM1, the 1995 Normative Addendum 1.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

ISO C99 ISO/IEC 9899:1999, the 1999 International Standard for C.

ISO C++ ISO/IEC 14822:2003, the 2003 International Standard for C++.

Errors

When `--strict` is in force and a violation of the relevant ISO standard occurs, the compiler issues an error message.

The severity of diagnostic messages can be controlled in the usual way.

Example

```
void foo(void)
{
    long long i; /* okay in nonstrict C90 */
}
```

Compiling this code with `--strict` generates an error.

See also

- *--c90* on page 3-18
- *--c99* on page 3-18
- *--cpp* on page 3-20
- *--strict_warnings*
- *Dollar signs in identifiers* on page 4-11
- *Source language modes of the compiler* on page 2-3 in *Using the Compiler*.

3.1.143 `--strict_warnings`

Diagnostics that are errors in `--strict` mode are downgraded to warnings, where possible. It is sometimes not possible for the compiler to downgrade a strict error, for example, where it cannot construct a legitimate program to recover.

Errors

When `--strict_warnings` is in force and a violation of the relevant ISO standard occurs, the compiler normally issues a warning message.

The severity of diagnostic messages can be controlled in the usual way.

Note

In some cases, the compiler issues an error message instead of a warning when it detects something that is strictly illegal, and terminates the compilation. For example:

```
#ifdef $Super$
extern void $Super$$__aeabi_idiv0(void); /* intercept __aeabi_idiv0 */
#endif
```

Compiling this code with `--strict_warnings` generates an error if you do not use the `--dollar` option.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example

```
void foo(void)
{
    long long i; /* okay in nonstrict C90 */
}
```

Compiling this code with `--strict_warnings` generates a warning message.

Compilation continues, even though the expression **`long long`** is strictly illegal.

See also

- *Source language modes* on page 2-3
- *Dollar signs in identifiers* on page 4-11
- *--c90* on page 3-18
- *--c99* on page 3-18
- *--cpp* on page 3-20
- *--strict, --no_strict* on page 3-88.

3.1.144 `--sys_include`

This option removes the current place from the include search path.

Quoted include files are treated in a similar way to angle-bracketed include files, except that quoted include files are always searched for first in the directories specified by `-I`, and angle-bracketed include files are searched for first in the `-J` directories.

See also

- *-Idir[,dir,...]* on page 3-49
- *-Jdir[,dir,...]* on page 3-55
- *--kandr_include* on page 3-55
- *--preinclude=filename* on page 3-78
- *Compiler search rules and the current place* on page 3-19 in *Using the ARM Compiler*
- *Compiler command-line options and search paths* on page 3-18 in *Using the ARM Compiler*.

3.1.145 `--thumb`

This option configures the compiler to target the Thumb instruction set.

Default

This is the default option for targets that do not support the ARM instruction set.

See also

- *--arm* on page 3-11
- *#pragma arm* on page 5-47
- *#pragma thumb* on page 5-60
- *ARM architectures supported by the toolchain* on page 2-14 in *Introducing the ARM Compiler toolchain*
- *Selecting the target CPU at compile time on page 5-8* in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.146 `--trigraphs`, `--no_trigraphs`

This option enables and disables trigraph recognition.

Default

The default is `--trigraphs`.

See also

- *ISO/IEC 9899:TC2*.

3.1.147 `--type_traits_helpers`, `--no_type_traits_helpers`

These options enable and disable support for C++ type traits helpers (such as `__is_union` and `__has_virtual_destructor`). Type traits helpers are enabled by default.

3.1.148 `-U`*`name`*

This option removes any initial definition of the macro *`name`*.

The macro *`name`* can be either:

- a predefined macro
- a macro specified using the `-D` option.

Note

Not all compiler predefined macros can be undefined.

Syntax

`-U`*`name`*

Where:

`name` is the name of the macro to be undefined.

Usage

Specifying `-U`*`name`* has the same effect as placing the text `#undef` *`name`* at the head of each source file.

Restrictions

The compiler defines and undefines macros in the following order:

1. compiler predefined macros
2. macros defined explicitly, using `-D`*`name`*
3. macros explicitly undefined, using `-U`*`name`*.

See also

- *-C* on page 3-17
- *-Dname[(parm-list)][=def]* on page 3-22
- *-E* on page 3-35
- *-M* on page 3-64
- *Compiler predefines* on page 5-98.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.149 `--unaligned_access`, `--no_unaligned_access`

These options enable and disable unaligned accesses to data on ARM architecture-based processors.

Default

The default is `--unaligned_access` on ARM-architecture based processors that support unaligned accesses to data. This includes:

- all ARMv6 architecture-based processors
- ARMv7-A, ARMv7-R, and ARMv7-M architecture-based processors.

The default is `--no_unaligned_access` on ARM-architecture based processors that do not support unaligned accesses to data. This includes:

- all pre-ARMv6 architecture-based processors
- ARMv6-M architecture-based processors.

Usage

`--unaligned_access`

Use `--unaligned_access` on processors that support unaligned accesses to data to speed up accesses to packed structures.

To enable unaligned support, you must:

- Clear the `A` bit, bit 1, of CP15 register 1 in your initialization code.
- Set the `U` bit, bit 22, of CP15 register 1 in your initialization code.

  The initial value of the `U` bit is determined by the **UBITINIT** input to the core. The MMU must be on, and the memory marked as normal memory.

The libraries include special versions of certain library functions designed to exploit unaligned accesses. When unaligned access support is enabled, the compilation tools use these library functions to take advantage of unaligned accesses.

`--no_unaligned_access`

Use `--no_unaligned_access` to disable the generation of unaligned word and halfword accesses on ARMv6 processors.

To enable modulo four-byte alignment checking on an ARMv6 target without unaligned accesses, you must:

- Set the `A` bit, bit 1, of CP15 register 1 in your initialization code.
- Set the `U` bit, bit 22, of CP15 register 1 in your initialization code.

  The initial value of the `U` bit is determined by the **UBITINIT** input to the core.

Note

Unaligned doubleword accesses, for example unaligned accesses to **`long long`** integers, are not supported by ARM processor cores. Doubleword accesses must be either eight-byte or four-byte aligned.

The compiler does not provide support for modulo eight-byte alignment checking. That is, the configuration `U = 0, A = 1` in CP15 register 1 is not supported by the compiler, or more generally, by the ARM compiler toolset.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

The libraries include special versions of certain library functions designed to exploit unaligned accesses. To prevent these enhanced library functions being used when unaligned access support is disabled, you have to specify `--no_unaligned_access` on both the compiler command line and the assembler command line when compiling a mixture of C and C++ source files and asssembly language source files.

Restrictions

Code compiled for processors supporting unaligned accesses to data can run correctly only if the choice of alignment support in software matches the choice of alignment support on the processor core.

See also

- *--cpu=name* on page 3-20
- *Assembler command line syntax* on page 2-2 *in the Assembler Reference*

3.1.150 `--use_frame_pointer`

This option enables `armcc` to store frame pointers in the stack in ARM and Thumb2 code. Using this option reserves `R11` to store the frame pointer. Do not use `R11` in inline assembly code when using this option.

See also

- *ARM registers* on page 3-8 in Using the Assembler
- *General-purpose registers* on page 3-10 in Using the Assembler.

3.1.151 `--use_pch=`*`filename`*

This option instructs the compiler to use a PCH file with the specified filename as part of the current compilation.

This option takes precedence if you include `--pch` on the same command line.

Syntax

`--use_pch=`*`filename`*

Where:

`filename` is the PCH file to be used as part of the current compilation.

Restrictions

The effect of this option is negated if you include `--create_pch=`*`filename`* on the same command line.

Errors

If the specified file does not exist, or is not a valid PCH file, the compiler generates an error.

See also

- *--create_pch=filename* on page 3-22
- *--pch* on page 3-75

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *--pch_dir=dir* on page 3-75
- *--pch_messages, --no_pch_messages* on page 3-76
- *--pch_verbose, --no_pch_verbose* on page 3-76
- *PreCompiled Header (PCH) files* on page 4-29 in *Using the Compiler*.

3.1.152 `--using_std`, `--no_using_std`

This option enables or disables implicit use of the `std` namespace when standard header files are included in C++.

Note

This option is provided only as a migration aid for legacy source code that does not conform to the C++ standard. Its use is not recommended.

Mode

This option is effective only if the source language is C++.

Default

The default is `--no_using_std`.

See also

- *Namespaces* on page 6-13.

3.1.153 `--version_number`

This option displays the version of `armcc` being used.

Example

```
armcc --version_number
```

The compiler prints the version number, for example, `400400`.

See also

- *--help* on page 3-49
- *--vsn* on page 3-96.

3.1.154 `--vfe`, `--no_vfe`

This option enables or disables *Virtual Function Elimination* (VFE) in C++.

VFE enables unused virtual functions to be removed from code. When VFE is enabled, the compiler places the information in special sections with the prefix `.arm_vfe_`. These sections are ignored by linkers that are not VFE-aware, because they are not referenced by the rest of the code. Therefore, they do not increase the size of the executable. However, they increase the size of the object files.

Mode

This option is effective only if the source language is C++.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Default

The default is `--vfe`, except for the case where legacy object files compiled with a pre-RVCT v2.1 compiler do not contain VFE information.

See also

- *Calling a pure virtual function* on page E-3
- *Elimination of unused virtual functions* on page 5-5 in *Using the Linker*.

3.1.155 `--via=filename`

This option instructs the compiler to read additional command-line options from a specified file. The options read from the file are added to the current command line.

Via commands can be nested within via files.

Syntax

`--via=filename`

Where:

filename is the name of a via file containing options to be included on the command line.

If *filename* is unspecified, the compiler faults use of `--via`.

Example

Given a source file `main.c`, a via file `apcs.txt` containing the line:

```
--apcs=/rwpi --no_lower_rwpi --via=L_apcs.txt
```

and a second via file `L_apcs.txt` containing the line:

```
-L--rwpi -L--callgraph
```

compiling `main.c` with the command line:

```
armcc main.c -L-o"main.axf" --via=apcs.txt
```

compiles `main.c` using the command line:

```
armcc --no_lower_rwpi --apcs=/rwpi -L--rwpi -L--callgraph -L-o"main.axf" main.c
```

See also

- Appendix B *Via File Syntax*
- *Using a text file to specify command-line options* on page 2-20 in *Introducing ARM Compilation Tools*.

3.1.156 `--visibility_inlines_hidden`

This option stops inline member functions acquiring dynamic linkage (`default` visibility) from:

- `class __declspec(dllexport)`
- a class visibility attribute
- `--no_hide_all`.

Non-member functions are not affected.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *__declspec(dllexport)* on page 5-19
- *--hide_all, --no_hide_all* on page 3-49.

3.1.157 `--vla`, `--no_vla`

This option enables or disables support for variable length arrays.

Default

C90 and Standard C++ do not support variable length arrays by default. Select the option `--vla` to enable support for variable length arrays in C90 or Standard C++.

Variable length arrays are supported in Standard C. The option `--vla` is implicitly selected either when the source language is C99.

Example

```
size_t arr_size(int n)
{
    char array[n];          // variable length array, dynamically allocated
    return sizeof array;   // evaluated at runtime
}
```

See also

- *--c90* on page 3-18
- *--c99* on page 3-18
- *--cpp* on page 3-20

3.1.158 `--vsn`

This option displays the version information and the license details.

See also

- *--help* on page 3-49
- *--version_number* on page 3-94.

3.1.159 `-W`

This option instructs the compiler to suppress all warning messages.

See also

- *--brief_diagnostics, --no_brief_diagnostics* on page 3-15
- *--diag_error=tag[,tag,...]* on page 3-30
- *--diag_remark=tag[,tag,...]* on page 3-31
- *--diag_style={arm|ide|gnu}* on page 3-31
- *--diag_suppress=tag[,tag,...]* on page 3-32
- *--diag_warning=tag[,tag,...]* on page 3-33
- *--errors=filename* on page 3-36
- *--remarks* on page 3-82
- *--wrap_diagnostics, --no_wrap_diagnostics* on page 3-98

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

3.1.160 --wchar, --no_wchar

This option permits or forbids the use of `wchar_t`. It does not necessarily fault declarations, providing they are unused.

Usage

Use this option to create an object file that is independent of `wchar_t` size.

Restrictions

If `--no_wchar` is specified:

- `wchar_t` fields in structure declarations are faulted by the compiler, regardless of whether or not the structure is used
- `wchar_t` in a typedef is faulted by the compiler, regardless of whether or not the typedef is used.

Default

The default is `--wchar`.

See also

- *--wchar16*
- *--wchar32*.

3.1.161 --wchar16

This option changes the type of `wchar_t` to **`unsigned short`**.

Selecting this option modifies both the type of the defined type `wchar_t` in C and the type of the native type **`wchar_t`** in C++. It also affects the values of `WCHAR_MIN` and `WCHAR_MAX`.

Default

The compiler assumes `--wchar16` unless `--wchar32` is explicitly specified.

See also

- *--wchar, --no_wchar*
- *--wchar32*
- *Predefined macros* on page 5-98.

3.1.162 --wchar32

This option changes the type of `wchar_t` to **`unsigned int`**.

Selecting this option modifies both the type of the defined type `wchar_t` in C and the type of the native type **`wchar_t`** in C++. It also affects the values of `WCHAR_MIN` and `WCHAR_MAX`.

Default

The compiler assumes `--wchar16` unless `--wchar32` is explicitly specified.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--wchar, --no_wchar* on page 3-97
- *--wchar16* on page 3-97
- *Predefined macros* on page 5-98.

3.1.163 `--whole_program`

This option promises the compiler that the source files specified on the command line form the whole program. The compiler is then able to apply optimizations based on the knowledge that the source code visible to it is the complete set of source code for the program being compiled. Without this knowledge, the compiler is more conservative when applying optimizations to the code.

Usage

Use this option to gain maximum performance from a small program.

Restriction

Do not use this option if you do not have all of the source code to give to the compiler.

See also

- *--multifile, --no_multifile* on page 3-67.

3.1.164 `--wrap_diagnostics`, `--no_wrap_diagnostics`

This option enables or disables the wrapping of error message text when it is too long to fit on a single line.

Default

The default is `--no_wrap_diagnostics`.

See also

- *--brief_diagnostics, --no_brief_diagnostics* on page 3-15
- *--diag_error=tag[,tag,...]* on page 3-30
- *--diag_remark=tag[,tag,...]* on page 3-31
- *--diag_style={arm|ide|gnu}* on page 3-31
- *--diag_suppress=tag[,tag,...]* on page 3-32
- *--diag_warning=tag[,tag,...]* on page 3-33
- *--errors=filename* on page 3-36
- *--remarks* on page 3-82
- *-W* on page 3-96
- Chapter 6 *Compiler Diagnostic Messages* in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Chapter 4
Language Extensions

This chapter describes the language extensions supported by the compiler, and includes:

- *Preprocessor extensions* on page 4-2
- *C99 language features available in C90* on page 4-4
- *C99 language features available in C++ and C90* on page 4-6
- *Standard C language extensions* on page 4-9
- *Standard C++ language extensions* on page 4-13
- *Standard C and Standard C++ language extensions* on page 4-16

For additional reference material on the compiler see also:

- Appendix D *Standard C Implementation Definition*
- Appendix E *Standard C++ Implementation Definition*
- Appendix F *C and C++ Compiler Implementation Limits*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

4.1 Preprocessor extensions

The compiler supports several extensions to the preprocessor, including the `#assert` preprocessing extensions of System V release 4.

4.1.1 #assert

The `#assert` preprocessing extensions of System V release 4 are permitted. These enable definition and testing of predicate names.

Such names are in a namespace distinct from all other names, including macro names.

Syntax

`#assert` *`name`*

`#assert` *`name`*`[(`*`token-sequence`*`)]`

Where:

`name` is a predicate name

`token-sequence` is an optional sequence of tokens.

If the token sequence is omitted, *`name`* is not given a value.

If the token sequence is included, *`name`* is given the value *`token-sequence`*.

Example

A predicate name defined using `#assert` can be tested in a `#if` expression, for example:

`#if #`*`name`*`(`*`token-sequence`*`)`

This has the value 1 if a `#assert` of the name *`name`* with the token-sequence *`token-sequence`* has appeared, and 0 otherwise. A given predicate can be given more than one value at a given time.

See also

- *#unassert*.

4.1.2 #unassert

A predicate name can be deleted using a `#unassert` preprocessing directive.

Syntax

`#unassert` *`name`*

`#unassert` *`name`*`[(`*`token-sequence`*`)]`

Where:

`name` is a predicate name

`token-sequence` is an optional sequence of tokens.

If the token sequence is omitted, all definitions of *`name`* are removed.

If the token sequence is included, only the indicated definition is removed. All other definitions are left intact.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *#assert* on page 4-2.

4.1.3 #warning

The preprocessing directive `#warning` is supported. Like the `#error` directive, this produces a user-defined warning at compilation time. However, it does not halt compilation.

Restrictions

The `#warning` directive is not available if the `--strict` option is specified. If used, it produces an error.

See also

- *--strict, --no_strict* on page 3-88.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

4.2 C99 language features available in C90

The compiler supports numerous extensions to the ISO C90 standard, for example, C99-style // comments.

These extensions are available if the source language is C90 and you are compiling in nonstrict mode.

These extensions are not available if the source language is C90 and the compiler is restricted to compiling strict C90 using the `--strict` compiler option.

Note

Language features of Standard C and Standard C++, for example C++-style // comments, might be similar to the C90 language extensions described in this section. Such features continue to remain available if you are compiling strict Standard C or strict Standard C++ using the `--strict` compiler option.

4.2.1 // comments

The character sequence // starts a one line comment, like in C99 or C++.

// comments in C90 have the same semantics as // comments in C99.

Example

```
// this is a comment
```

See also

- *New language features of C99 on page 5-77* in *Using the Compiler*.

4.2.2 Subscripting struct

In C90, arrays that are not lvalues still decay to pointers, and can be subscripted. However, you must not modify or use them after the next sequence point, and you must not apply the unary `&` operator to them. Arrays of this kind can be subscripted in C90, but they do not decay to pointers outside C99 mode.

Example

```
struct Subscripting_Struct
{
    int a[4];
};
extern struct Subscripting_Struct Subscripting_0(void);
int Subscripting_1 (int index)
{
    return Subscripting_0().a[index];
}
```

4.2.3 Flexible array members

The last member of a **`struct`** can have an incomplete array type. The last member must not be the only member of the **`struct`**, otherwise the **`struct`** is zero in size.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example

```
typedef struct
{
    int len;
    char p[]; // incomplete array type, for use in a malloc'd data structure
} str;
```

See also

- *New language features of C99 on page 5-77* in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

4.3 C99 language features available in C++ and C90

The compiler supports numerous extensions to the ISO C++ standard and to the C90 language, for example, function prototypes that override old-style nonprototype definitions.

These extensions are available if:

- the source language is C++ and you are compiling in nonstrict mode
- the source language is C90 and you are compiling in nonstrict mode.

These extensions are not available if:

- the source language is C++ and the compiler is restricted to compiling strict Standard C++ using the `--strict` compiler option.
- the source language is C90 and the compiler is restricted to compiling strict Standard C using the `--strict` compiler option.

Note

Language features of Standard C, for example **`long long`** integers, might be similar to the language extensions described in this section. Such features continue to remain available if you are compiling strict Standard C++ or strict C90 using the `--strict` compiler option.

4.3.1 Variadic macros

In C90 and C++ you can declare a macro to accept a variable number of arguments.

The syntax for declaring a variadic macro in C90 and C++ follows the C99 syntax for declaring a variadic macro.

Example

```
#define debug(format, ...) fprintf (stderr, format, __VA_ARGS__)
void variadic_macros(void)
{
    debug ("a test string is printed out along with %x %x %x\n", 12, 14, 20);
}
```

See also

- *New language features of C99 on page 5-77* in *Using the Compiler*.

4.3.2 long long

The ARM compiler supports 64-bit integer types through the type specifiers **`long long`** and **`unsigned long long`**. They behave analogously to **`long`** and **`unsigned long`** with respect to the usual arithmetic conversions. `__int64` is a synonym for **`long long`**.

Integer constants can have:

- an `ll` suffix to force the type of the constant to **`long long`**, if it fits, or to **`unsigned long long`** if it does not fit
- a `ull` or `llu` suffix to force the type of the constant to **`unsigned long long`**.

Format specifiers for `printf()` and `scanf()` can include `ll` to specify that the following conversion applies to a **`long long`** argument, as in `%lld` or `%llu`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Also, a plain integer constant is of type **long long** or **unsigned long long** if its value is large enough. There is a warning message from the compiler indicating the change. For example, in strict 1990 ISO Standard C 2147483648 has type **unsigned long**. In ARM C and C++ it has the type **long long**. One consequence of this is the value of an expression such as:

```
2147483648 > -1
```

This expression evaluates to 0 in strict C and C++, and to 1 in ARM C and C++.

The **long long** types are accommodated in the usual arithmetic conversions.

See also

- *__int64* on page 5-7.

4.3.3 restrict

The **restrict** keyword is a C99 feature. It enables you to convey a declaration of intent to the compiler that different pointers and function parameter arrays do not point to overlapping regions of memory at runtime. This enables the compiler to perform optimizations that can otherwise be prevented because of possible aliasing.

Usage

The keywords `__restrict` and `__restrict__` are supported as synonyms for **restrict** and are always available.

You can specify `--restrict` to allow the use of the **restrict** keyword in C90 or C++.

Restrictions

The declaration of intent is effectively a promise to the compiler that, if broken, results in undefined behavior.

Example

The following example shows use of the **restrict** keyword applied to function parameter arrays.

```
void copy_array(int n, int *restrict a, int *restrict b)
{
    while (n-- > 0)
        *a++ = *b++;
}
```

The following example shows use of the **restrict** keyword applied to different pointers that exist in the form of local variables.

```
void copy_bytes(int n, int *a, int *b)
{
  int *restrict x;
  int *restrict y;

    x = a;
    y = b;

    while (n-- > 0)
     *q++ = *s++;
}
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--restrict, --no_restrict* on page 3-83
- *New language features of C99 on page 5-77* in *Using the Compiler*.

4.3.4 Hexadecimal floats

C90 and C++ support floating-point numbers that can be written in hexadecimal format.

Example

```
float hex_floats(void)
{
    return 0x1.fp3;    // 1.55e1
}
```

See also

- *New language features of C99 on page 5-77* in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

4.4 Standard C language extensions

The compiler supports numerous extensions to the ISO C99 standard, for example, function prototypes that override old-style nonprototype definitions.

These extensions are available if:

- the source language is C99 and you are compiling in nonstrict mode
- the source language is C90 and you are compiling in nonstrict mode.

None of these extensions is available if:

- the source language is C90 and the compiler is restricted to compiling strict C90 using the `--strict` compiler option.
- the source language is C99 and the compiler is restricted to compiling strict Standard C using the `--strict` compiler option.
- the source language is C++.

4.4.1 Constant expressions

Extended constant expressions are supported in initializers. The following examples show the compiler behavior for the default, `--strict_warnings`, and `--strict` compiler modes.

Example 1, assigning the address of variable

Your code might contain constant expressions that assign the address of a variable at file scope, for example:

```
int i;
int j = (int)&i; /* but not allowed by ISO */
```

When compiling for C, this produces the following behavior:

- In default mode a warning is produced.
- In `--strict_warnings` mode a warning is produced.
- In `--strict` mode, an error is produced.

Example 2, constant value initializers

Table 4-1 compares the behavior of the ARM compilation tools with the ISO C Standard.

If compiling with `--strict_warnings` in place of `--strict`, the example source code that is not valid with `--strict` become valid. The `--strict` error message is downgraded to a warning message.

Table 4-1 Behavior of constant value initializers in comparison with ISO Standard C

Example source code	ISO C Standard	ARM compilation tools	
		`--strict` **mode**	**Nonstrict mode**
`extern int const c = 10;`	Valid	Valid	Valid
`extern int const x = c + 10;`	Not valid	Not valid	Valid
`static int y = c + 10;`	Not valid	Not valid	Valid
`static int const z = c + 10;`	Not valid	Not valid	Valid
`extern int *const cp = (int*)0x100;`	Valid	Valid	Valid

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table 4-1 Behavior of constant value initializers in comparison with ISO Standard C (continued)

Example source code	ISO C Standard	ARM compilation tools	
		`--strict` **mode**	**Nonstrict mode**
`extern int *const xp = cp + 0x100;`	Not valid	Not valid	Valid
`static int *yp = cp + 0x100;`	Not valid	Not valid	Valid
`static int *const zp = cp + 0x100;`	Not valid	Not valid	Valid

See also

- *--extended_initializers, --no_extended_initializers* on page 3-38
- *--strict, --no_strict* on page 3-88
- *--strict_warnings* on page 3-89.

4.4.2 Array and pointer extensions

The following array and pointer extensions are supported:

- Assignment and pointer differences are permitted between pointers to types that are interchangeable but not identical, for example, **`unsigned char *`** and **`char *`**. This includes pointers to same-sized integral types, typically, **`int *`** and **`long *`**. A warning is issued.

  Assignment of a string constant to a pointer to any kind of character is permitted without a warning.

- Assignment of pointer types is permitted in cases where the destination type has added type qualifiers that are not at the top level, for example, assigning **`int **`** to **`const int **`**. Comparisons and pointer difference of such pairs of pointer types are also permitted. A warning is issued.

- In operations on pointers, a pointer to `void` is always implicitly converted to another type if necessary. Also, a null pointer constant is always implicitly converted to a null pointer of the right type if necessary. In ISO C, some operators permit these, and others do not.

- Pointers to different function types can be assigned or compared for equality (`==`) or inequality (`!=`) without an explicit type cast. A warning or error is issued.

  This extension is prohibited in C++ mode.

- A pointer to **`void`** can be implicitly converted to, or from, a pointer to a function type.

- In an initializer, a pointer constant value can be cast to an integral type if the integral type is big enough to contain it.

- A non lvalue array expression is converted to a pointer to the first element of the array when it is subscripted or similarly used.

4.4.3 Block scope function declarations

Two extensions to block scope function declarations are supported:

- a block-scope function declaration also declares the function name at file scope
- a block-scope function declaration can have static storage class, thereby causing the resulting declaration to have internal linkage by default.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example

```
void f1(void)
{
    static void g(void); /* static function declared in local scope */
                         /* use of static keyword is illegal in strict ISO C */
}
void f2(void)
{
    g();                 /* uses previous local declaration */
}
static void g(int i)
{ } /* error - conflicts with previous declaration of g */
```

4.4.4 Dollar signs in identifiers

Dollar (`$`) signs are permitted in identifiers.

Note

When compiling with the `--strict` option, you can use the `--dollar` command-line option to permit dollar signs in identifiers.

Example

```
#define DOLLAR$
```

See also

- *--dollar, --no_dollar* on page 3-34
- *--strict, --no_strict* on page 3-88.

4.4.5 Top-level declarations

A C input file can contain no top-level declarations.

Errors

A remark is issued if a C input file contains no top-level declarations.

Note

Remarks are not displayed by default. To see remark messages, use the compiler option `--remarks`.

See also

- *--remarks* on page 3-82.

4.4.6 Benign redeclarations

Benign redeclarations of **typedef** names are permitted. That is, a **typedef** name can be redeclared in the same scope as the same type.

Example

```
typedef int INT;typedef int INT; /* redeclaration */
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

4.4.7 External entities

External entities declared in other scopes are visible.

Errors

The compiler generates a warning if an external entity declared in another scope is visible.

Example

```
void f1(void)
{
    extern void f();
}
void f2(void)
{
    f(); /* Out of scope declaration */
}
```

4.4.8 Function prototypes

The compiler recognizes function prototypes that override old-style nonprototype definitions that appear at a later position in your code.

Errors

The compiler generates a warning message if you use old-style function prototypes.

Example

```
int function_prototypes(char);
// Old-style function definition.
int function_prototypes(x)
    char x;
{
    return x == 0;
}
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

4.5 Standard C++ language extensions

The compiler supports numerous extensions to the ISO C++ standard, for example, qualified names in the declaration of class members.

These extensions are available if the source language is C++ and you are compiling in nonstrict mode.

These extensions are not available if the source language is C++ and the compiler is restricted to compiling strict Standard C++ using the `--strict` compiler option.

4.5.1 ? operator

A ? operator whose second and third operands are string literals or wide string literals can be implicitly converted to `char *` or `wchar_t *`. In C++ string literals are `const`. There is an implicit conversion that enables conversion of a string literal to `char *` or `wchar_t *`, dropping the `const`. That conversion, however, applies only to simple string literals. Permitting it for the result of a ? operation is an extension.

Example

```
char *p = x ? "abc" : "def";
```

4.5.2 Declaration of a class member

A qualified name can be used in the declaration of a class member.

Errors

A warning is issued if a qualified name is used in the declaration of a class member.

Example

```
struct A
{
    int A::f();  // is the same as int f();
};
```

4.5.3 `friend`

A **`friend`** declaration for a **`class`** can omit the class keyword.

Access checks are not carried out on **`friend`** declarations by default. Use the `--strict` command-line option to force access checking.

Example

```
class B;
class A
{
    friend B;  // is the same as "friend class B"
};
```

See also

- *--strict, --no_strict* on page 3-88.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

4.5.4 Read/write constants

A linkage specification for external constants indicates that a constant can be dynamically initialized or have mutable members.

Note

The use of "C++:read/write" linkage is only necessary for code compiled with --apcs /rwpi. If you recompile existing code with this option, you must change the linkage specification for external constants that are dynamically initialized or have mutable members.

Compiling C++ with the --apcs /rwpi option deviates from the ISO C++ Standard. The declarations in Example 4-1 assume that x is in a read-only segment.

Example 4-1 External access

```
extern const T x;
extern "C++" const T x;
extern "C" const T x;
```

Dynamic initialization of x including user-defined constructors is not possible for constants and T cannot contain mutable members. The new linkage specification in Example 4-2 declares that x is in a read/write segment even if it is initialized with a constant. Dynamic initialization of x is permitted and T can contain mutable members. The definitions of x, y, and z in another file must have the same linkage specifications.

Example 4-2 Linkage specification

```
extern const int z;                   /* in read-only segment, cannot  */
                                      /* be dynamically initialized    */
extern "C++:read/write" const int y;  /* in read/write segment */
                                      /* can be dynamically initialized */
extern "C++:read/write"
{
    const int i=5;                    /* placed in read-only segment, */
                                      /* not extern because implicitly static */
    extern const T x=6;               /* placed in read/write segment */
    struct S
    {
        static const T T x;           /* placed in read/write segment */
    };
}
```

Constant objects must not be redeclared with another linkage. The code in Example 4-3 produces a compile error.

Example 4-3 Compiler error

```
extern "C++"  const T x;
extern "C++:read/write"  const T x; /* error */
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Note

Because C does not have the linkage specifications, you cannot use a **`const`** object declared in C++ as `extern "C++:read/write"` from C.

See also

- *--apcs=qualifer...qualifier* on page 3-7.

4.5.5 Scalar type constants

Constants of scalar type can be defined within classes. This is an old form. The modern form uses an initialized static data member.

Errors

A warning is issued if you define a member of constant integral type within a class.

Example

```
class A
{
    const int size = 10; // must be static const int size = 10;
    int a[size];
};
```

4.5.6 Specialization of nonmember function templates

As an extension, it is permitted to specify a storage class on a specialization of a nonmember function template.

4.5.7 Type conversions

Type conversion between a pointer to an `extern "C"` function and a pointer to an `extern "C++"` function is permitted.

Example

```
extern "C" void f();    // f's type has extern "C" linkage
void (*pf)() = &f;      // pf points to an extern "C++" function
                        // error unless implicit conversion is allowed
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

4.6 Standard C and Standard C++ language extensions

The compiler supports numerous extensions to both the ISO C99 and the ISO C++ Standards, such as various integral type extensions, various floating-point extensions, hexadecimal floating-point constants, and anonymous classes, structures, and unions.

These extensions are available if:

- the source language is C++ and you are compiling in nonstrict mode
- the source language is C99 and you are compiling in nonstrict mode
- the source language is C90 and you are compiling in nonstrict mode.

These extensions are not available if:

- the source language is C++ and the compiler is restricted to compiling strict C++ using the `--strict` compiler option.
- the source language is C99 and the compiler is restricted to compiling strict Standard C using the `--strict` compiler option.
- the source language is C90 and the compiler is restricted to compiling strict C90 using the `--strict` compiler option.

4.6.1 Address of a register variable

The address of a variable with **`register`** storage class can be taken.

Errors

The compiler generates a warning if you take the address of a variable with **`register`** storage class.

Example

```
void foo(void)
{
    register int i;
    int *j = &i;
}
```

4.6.2 Arguments to functions

Default arguments can be specified for function parameters other than those of a top-level function declaration. For example, they are accepted on `typedef` declarations and on pointer-to-function and pointer-to-member-function declarations.

4.6.3 Anonymous classes, structures and unions

Anonymous classes, structures, and unions are supported as an extension. Anonymous structures and unions are supported in C and C++.

Anonymous unions are available by default in C++. However, you must specify the `anon_unions` pragma if you want to use:

- anonymous unions and structures in C
- anonymous classes and structures in C++.

An anonymous union can be introduced into a containing class by a **`typedef`** name. Unlike a true anonymous union, it does not have to be declared directly. For example:

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

```
typedef union
{
    int i, j;
} U;                // U identifies a reusable anonymous union.
#pragma anon_unions
class A
{
    U;              // Okay -- references to A::i and A::j are allowed.
};
```

The extension also enables anonymous classes and anonymous structures, as long as they have no C++ features. For example, no static data members or member functions, no nonpublic members, and no nested types (except anonymous classes, structures, or unions) are allowed in anonymous classes and anonymous structures. For example:

```
#pragma anon_unions
struct A
{
    struct
    {
        int i, j;
    };              // Okay -- references to i and j
};                  // through class A are allowed.

int foo(int m)
{
    A a;
     a.i = m;
     return a.i;
}
```

See also

- *#pragma anon_unions, #pragma no_anon_unions* on page 5-47.

4.6.4 Assembler labels

Assembler labels specify the assembler name to use for a C symbol. For example, you might have assembler code and C code that uses the same symbol name, such as `counter`. Therefore, you can export a different name to be used by the assembler:

```
int counter __asm__("counter_v1") = 0;
```

This exports the symbol `counter_v1` and not the symbol `counter`.

See also

- *__asm* on page 5-4.

4.6.5 Empty declaration

An empty declaration, that is a semicolon with nothing before it, is permitted.

Example

```
; // do nothing
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

4.6.6 Hexadecimal floating-point constants

The ARM compiler implements an extension to the syntax of numeric constants in C to enable explicit specification of floating-point constants as IEEE bit patterns.

Syntax

The syntax for specifying floating-point constants as IEEE bit patterns is:

`0f_`*`n`* Interpret an 8-digit hex number *n* as a **`float`** constant. There must be exactly eight digits.

`0d_`*`nn`* Interpret a 16-digit hex number *nn* as a **`double`** constant. There must be exactly 16 digits.

4.6.7 Incomplete enums

Forward declarations of `enums` are supported.

Example

```
enum Incomplete_Enums_0;
int Incomplete_Enums_2 (enum Incomplete_Enums_0 * passon)
{
    return 0;
}
int Incomplete_Enums_1 (enum Incomplete_Enums_0 * passon)
{
    return Incomplete_Enums_2(passon);
}
enum Incomplete_Enums_0 { ALPHA, BETA, GAMMA };
```

4.6.8 Integral type extensions

In an integral constant expression, an integral constant can be cast to a pointer type and then back to an integral type.

4.6.9 Label definitions

In Standard C and Standard C++, a statement must follow a label definition. In C and C++, a label definition can be followed immediately by a right brace.

Errors

The compiler generates a warning if a label definition is followed immediately by a right brace.

Example

```
void foo(char *p)
{
    if (p)
    {
        /* ... */
label:
    }
}
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

4.6.10 Long float

long float is accepted as a synonym for **double**.

4.6.11 Nonstatic local variables

Nonstatic local variables of an enclosing function can be referenced in a non-evaluated expression, for example, a `sizeof` expression inside a local class. A warning is issued.

4.6.12 Structure, union, enum, and bitfield extensions

The following structure, union, enum, and bitfield extensions are supported:

- In C, the element type of a file-scope array can be an incomplete **struct** or **union** type. The element type must be completed before its size is needed, for example, if the array is subscripted. If the array is not **extern**, the element type must be completed by the end of the compilation.
- The final semicolon preceding the closing brace } of a **struct** or **union** specifier can be omitted. A warning is issued.
- An initializer expression that is a single value and is used to initialize an entire static array, **struct**, or **union**, does not have to be enclosed in braces. ISO C requires the braces.
- An extension is supported to enable constructs similar to C++ anonymous unions, including the following:
  - not only anonymous unions but also anonymous structs are permitted. The members of anonymous structs are promoted to the scope of the containing **struct** and looked up like ordinary members.
  - they can be introduced into the containing **struct** by a **typedef** name. That is, they do not have to be declared directly, as is the case with true anonymous unions.
  - a tag can be declared but only in C mode.

  To enable support for anonymous structures and unions, you must use the `anon_unions` pragma.
- An extra comma is permitted at the end of an **enum** list but a remark is issued.
- **enum** tags can be incomplete. You can define the tag name and resolve it later, by specifying the brace-enclosed list.
- The values of enumeration constants can be given by expressions that evaluate to unsigned quantities that fit in the **unsigned int** range but not in the **int** range. For example:

  ```
  /* When ints are 32 bits: */
  enum a { w = -2147483648 };  /* No error */
  enum b { x = 0x80000000 };   /* No error */
  enum c { y = 0x80000001 };   /* No error */
  enum d { z = 2147483649 };   /* Error */
  ```
- Bit fields can have base types that are **enum** types or integral types besides **int** and **unsigned int**.

See also

- *Pragmas* on page 5-47
- *Structure, union, enum, and bitfield extensions*
- *New language features of C99 on page 5-77* in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Chapter 5
Compiler-specific Features

This chapter describes the compiler-specific features, and includes:

- *Keywords and operators* on page 5-2
- *__declspec attributes* on page 5-19
- *Function attributes* on page 5-25
- *Type attributes* on page 5-36
- *Variable attributes* on page 5-39
- *Pragmas* on page 5-47
- *Instruction intrinsics* on page 5-61
- *ARMv6 SIMD intrinsics* on page 5-88
- *ETSI basic operations* on page 5-89
- *C55x intrinsics* on page 5-91
- *VFP status intrinsic* on page 5-92
- *Fused Multiply Add (FMA) intrinsics* on page 5-93
- *Named register variables* on page 5-94
- *Compiler predefines* on page 5-98.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.1 Keywords and operators

This section describes the function keywords and operators supported by the compiler `armcc`.

Table 5-1 lists keywords that are ARM extensions to the C and C++ Standards. Standard C and Standard C++ keywords that do not have behavior or restrictions specific to the ARM compiler are not documented in the table.

Table 5-1 Keyword extensions supported by the ARM compiler

Keywords		
`__align`	`__int64`	`__svc`
`__ALIGNOF__`	`__INTADDR__`	`__svc_indirect`
`__asm`	`__irq`	`__svc_indirect_r7`
`__declspec`	`__packed`	`__value_in_regs`
`__forceinline`	`__pure`	`__weak`
`__global_reg`	`__softfp`	`__writeonly`
`__inline`	`__smc`	

5.1.1 `__align`

The `__align` keyword instructs the compiler to align a variable on an *n*-byte boundary.

`__align` is a storage class modifier. It does not affect the type of the function.

Syntax

`__align(n)`

Where:

n is the alignment boundary.

For local variables, *n* can take the values 1, 2, 4, or 8.

For global variables, *n* can take any value up to 0x80000000 in powers of 2.

Usage

`__align(n)` is useful when the normal alignment of the variable being declared is less than *n*. Eight-byte alignment can give a significant performance advantage with VFP instructions.

`__align` can be used in conjunction with **`extern`** and **`static`**.

Restrictions

Because `__align` is a storage class modifier, it cannot be used on:

- types, including **`typedef`**s and structure definitions
- function parameters.

You can only overalign. That is, you can make a two-byte object four-byte aligned but you cannot align a four-byte object at 2 bytes.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Examples

```
__align(8) char buffer[128];  // buffer starts on eight-byte boundary

void foo(void)
{
    ...
    __align(16) int i; // this alignment value is not permitted for
                       // a local variable
    ...
}

__align(16) int i; // permitted as a global variable.
```

See also

- *--min_array_alignment=opt* on page 3-66 in *Using the Compiler*.

5.1.2 __alignof__

The `__alignof__` keyword enables you to enquire about the alignment of a type or variable.

Note

This keyword is a GNU compiler extension that is supported by the ARM compiler.

Syntax

`__alignof__(type)`

`__alignof__(expr)`

Where:

type is a type

expr is an lvalue.

Return value

`__alignof__(type)` returns the alignment requirement for the type *type*, or 1 if there is no alignment requirement.

`__alignof__(expr)` returns the alignment requirement for the type of the lvalue *expr*, or 1 if there is no alignment requirement.

Example

```
int Alignment_0(void)
{
    return __alignof__(int);
}
```

See also

- *__ALIGNOF__*.

5.1.3 __ALIGNOF__

The `__ALIGNOF__` keyword returns the alignment requirement for a specified type, or for the type of a specified object.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Syntax

`__ALIGNOF__(`*`type`*`)`

`__ALIGNOF__(`*`expr`*`)`

Where:

`type` is a type

`expr` is an lvalue.

Return value

`__ALIGNOF__(`*`type`*`)` returns the alignment requirement for the type *`type`*, or 1 if there is no alignment requirement.

`__ALIGNOF__(`*`expr`*`)` returns the alignment requirement for the type of the lvalue *`expr`*, or 1 if there is no alignment requirement. The lvalue itself is not evaluated.

Example

```
typedef struct s_foo { int i; short j; } foo;
typedef __packed struct s_bar { int i; short j; } bar;
return __ALIGNOF(struct s_foo); // returns 4
return __ALIGNOF(foo);          // returns 4
return __ALIGNOF(bar);          // returns 1
```

See also

- *__alignof__* on page 5-3.

5.1.4 `__asm`

This keyword is used to pass information from the compiler to the ARM assembler `armasm`.

The precise action of this keyword depends on its usage.

Usage

Embedded assembler

The `__asm` keyword can be used to declare or define an embedded assembly function. For example:

```
__asm void my_strcpy(const char *src, char *dst);
```

Compiler support for embedded assembler on page 7-36 in *Using the Compiler* for more information.

Inline assembler

The `__asm` keyword can be used to incorporate inline assembly into a function. For example:

```
int qadd(int i, int j)
{
    int res;
    __asm
    {
        QADD   res, i, j
    }
    return res;
}
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See *Compiler support for inline assembly language* on page 7-4 in *Using the Compiler* for more information.

Assembler labels

The `__asm` keyword can be used to specify an assembler label for a C symbol. For example:

```
int count __asm__("count_v1"); // export count_v1, not count
```

See *Assembler labels* on page 4-17 for more information.

Named register variables

The `__asm` keyword can be used to declare a named register variable. For example:

```
register int foo __asm("r0");
```

See *Named register variables* on page 5-94 for more information.

5.1.5 __forceinline

The `__forceinline` keyword forces the compiler to compile a C or C++ function inline.

The semantics of `__forceinline` are exactly the same as those of the C++ **`inline`** keyword. The compiler attempts to inline a function qualified as `__forceinline`, regardless of its characteristics. However, the compiler does not inline a function if doing so causes problems. For example, a recursive function is inlined into itself only once.

`__forceinline` is a storage class qualifier. It does not affect the type of a function.

Note

This keyword has the function attribute equivalent `__attribute__((always_inline))`.

Example

```
__forceinline static int max(int x, in y)
{
    return x > y ? x : y; // always inline if possible
}
```

See also

- *--forceinline* on page 3-40
- *__attribute__((always_inline)) function attribute* on page 5-26.

5.1.6 __global_reg

The `__global_reg` storage class specifier allocates the declared variable to a global variable register.

Syntax

`__global_reg(`*`n`*`)` *`type varName`*

Where:

n Is an integer between one and eight.

type Is one of the following types:

- any integer type, except **`long long`**

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- any char type
- any pointer type.

`varName` Is the name of a variable.

Restrictions

If you use this storage class, you cannot use any additional storage class such as **`extern`**, **`static`**, or **`typedef`**.

In C, global register variables cannot be qualified or initialized at declaration. In C++, any initialization is treated as a dynamic initialization.

The number of available registers varies depending on the variant of the AAPCS being used, there are between five and seven registers available for use as global variable registers.

In practice, it is recommended that you do not use more than:

- three global register variables in ARM or Thumb-2
- one global register variable in Thumb-1
- half the number of available floating-point registers as global floating-point register variables.

If you declare too many global variables, code size increases significantly. In some cases, your program might not compile.

Caution

You must take care when using global register variables because:

- There is no check at link time to ensure that direct calls between different compilation units are sensible. If possible, define global register variables used in a program in each compilation unit of the program. In general, it is best to place the definition in a global header file. You must set up the value in the global register early in your code, before the register is used.
- A global register variable maps to a callee-saved register, so its value is saved and restored across a call to a function in a compilation unit that does not use it as a global register variable, such as a library function.
- Calls back into a compilation unit that uses a global register variable are dangerous. For example, if a function using a global register is called from a compilation unit that does not declare the global register variable, the function reads the wrong values from its supposed global register variables.
- This storage class can only be used at file scope.

Example

Example 5-1 declares a global variable register allocated to `r5`.

Example 5-1 Declaring a global integer register variable

```
__global_reg(2) int x; v2 is the synonym for r5
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example 5-2 produces an error because global registers must be specified in all declarations of the same variable.

Example 5-2 Global register - declaration error

```
int x;
__global_reg(1) int x; // error
```

In C, `__global_reg` variables cannot be initialized at definition. Example 5-3 produces an error in C, but not in C++.

Example 5-3 Global register - initialization error

```
__global_reg(1) int x=1; // error in C, OK in C++
```

See also

- *--global_reg=reg_name[,reg_name,...]* on page 3-47.

5.1.7 `__inline`

The `__inline` keyword suggests to the compiler that it compiles a C or C++ function inline, if it is sensible to do so.

The semantics of `__inline` are exactly the same as those of the **`inline`** keyword. However, **`inline`** is not available in C90.

`__inline` is a storage class qualifier. It does not affect the type of a function.

Example

```
__inline int f(int x)
{
    return x*5+1;
}
int g(int x, int y)
{
    return f(x) + f(y);
}
```

See also

- *Inline functions on page 5-29* in *Using the ARM Compiler*.

5.1.8 `__int64`

The `__int64` keyword is a synonym for the keyword sequence **`long long`**.

`__int64` is accepted even when using `--strict`.

See also

- *--strict, --no_strict* on page 3-88
- *long long* on page 4-6.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.1.9 `__INTADDR__`

The `__INTADDR__` operation treats the enclosed expression as a constant expression, and converts it to an integer constant.

Note

This is used in the `offsetof` macro.

Syntax

`__INTADDR(`*`expr`*`)`

Where:

`expr` is an integral constant expression.

Return value

`__INTADDR__(`*`expr`*`)` returns an integer constant equivalent to *`expr`*.

See also

- *Restrictions on embedded assembly language functions in C and C++ code* on page 7-40 in *Using the Compiler*.

5.1.10 `__irq`

The `__irq` keyword enables a C or C++ function to be used as an interrupt routine.

`__irq` is a function qualifier. It affects the type of the function.

Restrictions

All corrupted registers except floating-point registers are preserved, not only those that are normally preserved under the AAPCS. The default AAPCS mode must be used.

No arguments or return values can be used with `__irq` functions.

In architectures other than ARMv6-M and ARMv7-M, the function exits by setting the program counter to `lr`-4 and the `CPSR` to the value in `SPSR`.

Note

In ARMv6-M and ARMv7-M, specifying `__irq` does not affect the behavior of the compiled output. However, ARM recommends using `__irq` on exception handlers for clarity and easier software porting.

Note

- For architectures that support ARM and 32-bit Thumb, for example ARMv6T2, ARMv7-A, and ARMv7-R, functions specified as __irq compile to ARM or Thumb code depending on whether the compile option or `#pragma` specify ARM or Thumb.
- For Thumb only architectures, for example ARMv6-M and ARMv7-M, functions specified as `__irq` compile to Thumb code.
- For architectures before ARMv6T2, functions specified as `__irq` compile to ARM code even if you compile with `--thumb` or `#pragma thumb`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--thumb* on page 3-90
- *--arm* on page 3-11
- *#pragma thumb* on page 5-60
- *#pragma arm* on page 5-47
- *ARM, Thumb, and ThumbEE instruction sets* on page 3-3 in Using the Assembler

5.1.11 `__packed`

The `__packed` qualifier sets the alignment of any valid type to 1. This means that:

- there is no padding inserted to align the packed object
- objects of packed type are read or written using unaligned accesses.

The `__packed` qualifier applies to all members of a structure or union when it is declared using `__packed`. There is no padding between members, or at the end of the structure. All substructures of a packed structure must be declared using `__packed`. Integral subfields of an unpacked structure can be packed individually.

Usage

The `__packed` qualifier is useful to map a structure to an external data structure, or for accessing unaligned data, but it is generally not useful to save data size because of the relatively high cost of unaligned access. Only packing fields in a structure that requires packing can reduce the number of unaligned accesses.

Note

On ARM processors that do not support unaligned access in hardware, for example, pre-ARMv6, access to unaligned data can be costly in terms of code size and execution speed. Data accesses through packed structures must be minimized to avoid increase in code size and performance loss.

Restrictions

The following restrictions apply to the use of `__packed`:

- The `__packed` qualifier cannot be used on structures that were previously declared without `__packed`.
- Unlike other type qualifiers you cannot have both a `__packed` and non-`__packed` version of the same structure type.
- The `__packed` qualifier does not affect local variables of integral type.
- A packed structure or union is not assignment-compatible with the corresponding unpacked structure. Because the structures have a different memory layout, the only way to assign a packed structure to an unpacked structure is by a field-by-field copy.
- The effect of casting away `__packed` is undefined, except on **`char`** types. The effect of casting a nonpacked structure to a packed structure, or a packed structure to a nonpacked structure, is undefined. A pointer to an integral type that is not packed can be legally cast, explicitly or implicitly, to a pointer to a packed integral type.
- There are no packed array types. A packed array is an array of objects of packed type. There is no padding in the array.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example

Example 5-4 shows that a pointer can point to a packed type.

Example 5-4 Pointer to packed

```
typedef __packed int* PpI;          /* pointer to a __packed int */
__packed int *p;                    /* pointer to a __packed int */
PpI p2;                             /* 'p2' has the same type as 'p' */
                                    /* __packed is a qualifier  */
                                    /* like 'const' or 'volatile' */
typedef int *PI;                    /* pointer to int */
__packed PI p3;                     /* a __packed pointer to a normal int */
                                    /*  -- not the same type as 'p' and 'p2' */
int *__packed p4;                   /* 'p4' has the same type as 'p3' */
```

Example 5-5 shows that when a packed object is accessed using a pointer, the compiler generates code that works and that is independent of the pointer alignment.

Example 5-5 Packed structure

```
typedef __packed struct
{
    char x;                     // all fields inherit the __packed qualifier
    int y;
} X;                            // 5 byte structure, natural alignment = 1
int f(X *p)
{
    return p->y;                // does an unaligned read
}
typedef struct
{
    short x;
    char y;
    __packed int z;             // only pack this field
    char a;
} Y;                            // 8 byte structure, natural alignment = 2
int g(Y *p)
{
    return p->z + p->x;         // only unaligned read for z
}
```

See also

- *__attribute__((packed)) variable attribute* on page 5-42
- *#pragma pack(n)* on page 5-56
- *Packed structures* on page 6-8
- *The __packed qualifier and unaligned data access in C and C++ code on page 5-46* in *Using the Compiler*
- *Detailed comparison of an unpacked struct, a __packed struct, and a struct with individually __packed fields on page 5-51* in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.1.12 __pure

The `__pure` keyword asserts that a function declaration is pure.

A function is *pure* only if:

- the result depends exclusively on the values of its arguments
- the function has no side effects.

`__pure` is a function qualifier. It affects the type of a function.

Note

This keyword has the function attribute equivalent `__attribute__((const))`.

Default

By default, functions are assumed to be impure.

Usage

Pure functions are candidates for common subexpression elimination.

Restrictions

A function that is declared as pure can have no side effects. For example, pure functions:

- cannot call impure functions
- cannot use global variables or dereference pointers, because the compiler assumes that the function does not access memory, except stack memory
- must return the same value each time when called twice with the same parameters.

Example

```
int factr(int n) __pure
{
    int f = 1;
    while (n > 0)
        f *= n--;
    return f;}
```

See also

- *__attribute__((const)) function attribute* on page 5-27
- *Functions that return the same result when called with the same arguments on page 5-24* in *Using the Compiler*
- *Recommendation of postfix syntax when qualifying functions with ARM function modifiers on page 5-27* in *Using the Compiler*.

5.1.13 __smc

The `__smc` keyword declares an SMC (*Secure Monitor Call*) function. A call to the SMC function inserts an `SMC` instruction into the instruction stream generated by the compiler at the point of function invocation.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Note

The SMC instruction replaces the SMI instruction used in previous versions of the ARM assembly language.

__smc is a function qualifier. It affects the type of a function.

Syntax

```
__smc(int smc_num) return-type function-name([argument-list]);
```

Where:

smc_num Is a 4-bit immediate value used in the SMC instruction.

The value of *smc_num* is ignored by the ARM processor, but can be used by the SMC exception handler to determine what service is being requested.

Restrictions

The SMC instruction is available for selected ARM architecture-based processors, if they have the Security Extensions. See *SMC* on page 3-141 in the *Assembler Reference* for more information.

The compiler generates an error if you compile source code containing the __smc keyword for an architecture that does not support the SMC instruction.

Example

```
__smc(5) void mycall(void); /* declare a name by which SMC #5 can be called */
...
mycall();                   /* invoke the function */
```

See also

- *--cpu=name* on page 3-20
- *SMC* on page 3-141 in the *Assembler Reference*.

5.1.14 __softfp

The __softfp keyword asserts that a function uses software floating-point linkage.

__softfp is a function qualifier. It affects the type of the function.

Note

This keyword has the #pragma equivalent #pragma __softfp_linkage.

Usage

Calls to the function pass floating-point arguments in integer registers. If the result is a floating-point value, the value is returned in integer registers. This duplicates the behavior of compilation targeting software floating-point.

This keyword enables the same library to be used by sources compiled to use hardware and software floating-point.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Note

In C++, if a virtual function qualified with the `__softfp` keyword is to be overridden, the overriding function must also be declared as `__softfp`. If the functions do not match, the compiler generates an error.

See also

- *__attribute__((pcs("calling_convention")))* on page 5-31
- *--fpu=name* on page 3-44
- *#pragma softfp_linkage, #pragma no_softfp_linkage* on page 5-57
- *Compiler support for floating-point computations and linkage on page 5-63* in *Using the Compiler*.

5.1.15 __svc

The `__svc` keyword declares a *SuperVisor Call* (SVC) function taking up to four integer-like arguments and returning up to four results in a `value_in_regs` structure.

`__svc` is a function qualifier. It affects the type of a function.

Syntax

```
__svc(int svc_num) return-type function-name([argument-list]);
```

Where:

svc_num Is the immediate value used in the `SVC` instruction.

It is an expression evaluating to an integer in the range:

- 0 to 2^{24}–1 (a 24-bit value) in an ARM instruction
- 0-255 (an 8-bit value) in a 16-bit Thumb instruction.

Usage

This causes function invocations to be compiled inline as an AAPCS-compliant operation that behaves similarly to a normal call to a function.

The `__value_in_regs` qualifier can be used to specify that a small structure of up to 16 bytes is returned in registers, rather than by the usual structure-passing mechanism defined in the AAPCS.

Example

```
__svc(42) void terminate_1(int procnum); // terminate_1 returns no results
__svc(42) int terminate_2(int procnum);  // terminate_2 returns one result
typedef struct res_type
{
    int res_1;
    int res_2;
    int res_3;
    int res_4;
} res_type;
__svc(42) __value_in_regs res_type terminate_3(int procnum);
                                         // terminate_3 returns more than
                                         // one result
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Errors

When an ARM architecture variant or ARM architecture-based processor that does not support an `SVC` instruction is specified on the command line using the `--cpu` option, the compiler generates an error.

See also

- *--cpu=name* on page 3-20
- *__value_in_regs* on page 5-16
- *SVC* on page 3-135 in the *Assembler Reference*.

5.1.16 `__svc_indirect`

The `__svc_indirect` keyword passes an operation code to the SVC handler in `r12`.

`__svc_indirect` is a function qualifier. It affects the type of a function.

Syntax

```
__svc_indirect(int svc_num)
        return-type function-name(int real_num[, argument-list]);
```

Where:

`svc_num` Is the immediate value used in the `SVC` instruction.

It is an expression evaluating to an integer in the range:

- 0 to $2^{24}-1$ (a 24-bit value) in an ARM instruction
- 0-255 (an 8-bit value) in a 16-bit Thumb instruction.

`real_num` Is the value passed in `r12` to the handler to determine the function to perform.

To use the indirect mechanism, your system handlers must make use of the `r12` value to select the required operation.

Usage

You can use this feature to implement indirect SVCs.

Example

```
int __svc_indirect(0) ioctl(int svcino, int fn, void *argp);
```

Calling:

```
ioctl(IOCTL+4, RESET, NULL);
```

compiles to `SVC #0` with `IOCTL+4` in `r12`.

Errors

When an ARM architecture variant or ARM architecture-based processor that does not support an `SVC` instruction is specified on the command line using the `--cpu` option, the compiler generates an error.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--cpu=name* on page 3-20
- *__value_in_regs* on page 5-16
- *SVC* on page 3-135 in the *Assembler Reference*.

5.1.17 `__svc_indirect_r7`

The `__svc_indirect_r7` keyword behaves like `__svc_indirect`, but uses `r7` instead of `r12`.

`__svc_indirect_r7` is a function qualifier. It affects the type of a function.

Syntax

```
__svc_indirect_r7(int svc_num)
        return-type function-name(int real_num[, argument-list]);
```

Where:

`svc_num` Is the immediate value used in the `SVC` instruction.

It is an expression evaluating to an integer in the range:

- 0 to 2^{24}–1 (a 24-bit value) in an ARM instruction
- 0-255 (an 8-bit value) in a 16-bit Thumb instruction.

`real_num` Is the value passed in `r7` to the handler to determine the function to perform.

Usage

You can use this feature to implement indirect SVCs.

Example

```
long __svc_indirect_r7(0) \
        SVC_write(unsigned, int fd, const char * buf, size_t count);
#define write(fd, buf, count) SVC_write(4, (fd), (buf), (count))
```

Calling:

```
write(fd, buf, count);
```

compiles to `SVC #0` with `r0` = `fd`, `r1` = `buf`, `r2` = `count`, and `r7` = `4`.

Errors

When an ARM architecture variant or ARM architecture-based processor that does not support an `SVC` instruction is specified on the command line using the `--cpu` option, the compiler generates an error.

See also

- *__value_in_regs* on page 5-16
- *--cpu=name* on page 3-20
- *SVC* on page 3-135 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.1.18 `__value_in_regs`

The `__value_in_regs` qualifier instructs the compiler to return a structure of up to four integer words in integer registers or up to four floats or doubles in floating-point registers rather than using memory.

`__value_in_regs` is a function qualifier. It affects the type of a function.

Syntax

```
__value_in_regs return-type function-name([argument-list]);
```

Where:

`return-type` is the type of a structure of up to four words in size.

Usage

Declaring a function `__value_in_regs` can be useful when calling functions that return more than one result.

Restrictions

A C++ function cannot return a `__value_in_regs` structure if the structure requires copy constructing.

If a virtual function declared as `__value_in_regs` is to be overridden, the overriding function must also be declared as `__value_in_regs`. If the functions do not match, the compiler generates an error.

Errors

Where the structure returned in a function qualified by `__value_in_regs` is too big, a warning is produced and the `__value_in_regs` structure is then ignored.

Example

```
typedef struct int64_struct
{
    unsigned int lo;
    unsigned int hi;
} int64_struct;
__value_in_regs extern
    int64_struct mul64(unsigned a, unsigned b);
```

See also

- *Functions that return multiple values through registers on page 5-23* in *Using the ARM Compiler*.

5.1.19 `__weak`

This keyword instructs the compiler to export symbols weakly.

The `__weak` keyword can be applied to function and variable declarations, and to function definitions.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Usage

Functions and variable declarations

For declarations, this storage class specifies an **extern** object declaration that, even if not present, does not cause the linker to fault an unresolved reference.

For example:

```
__weak void f(void);
...
f(); // call f weakly
```

If the reference to a missing weak function is made from code that compiles to a branch or branch link instruction, then either:

- The reference is resolved as branching to the next instruction. This effectively makes the branch a `NOP`.
- The branch is replaced by a `NOP` instruction.

Function definitions

Functions defined with `__weak` export their symbols weakly. A weakly defined function behaves like a normally defined function unless a nonweakly defined function of the same name is linked into the same image. If both a nonweakly defined function and a weakly defined function exist in the same image then all calls to the function resolve to call the nonweak function. If multiple weak definitions are available, the linker chooses one for use by all calls.

Functions declared with `__weak` and then defined without `__weak` behave as nonweak functions.

Restrictions

There are restrictions when you qualify function and variable declarations, and function definitions, with `__weak`.

Functions and variable declarations

A function or variable cannot be used both weakly and nonweakly in the same compilation. For example, the following code uses `f()` weakly from `g()` and `h()`:

```
void f(void);
void g()
{
    f();
}
__weak void f(void);
void h()
{
    f();
}
```

It is not possible to use a function or variable weakly from the same compilation that defines the function or variable. The following code uses `f()` nonweakly from `h()`:

```
__weak void f(void);
void h()
{
    f();
}
void f() {}
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

The linker does not load the function or variable from a library unless another compilation uses the function or variable nonweakly. If the reference remains unresolved, its value is assumed to be `NULL`. Unresolved references, however, are not `NULL` if the reference is from code to a position-independent section or to a missing `__weak` function.

Function definitions

Weakly defined functions cannot be inlined.

Example

```
__weak const int c;           // assume 'c' is not present in final link
const int *f1() { return &c; } // '&c' returns non-NULL if
                               // compiled and linked /ropi
__weak int i;                 // assume 'i' is not present in final link
int *f2() { return &i; }      // '&i' returns non-NULL if
                               // compiled and linked /rwpi
__weak void f(void);          // assume 'f' is not present in final link
typedef void (*FP)(void);
FP g() { return f; }          // 'g' returns non-NULL if
                               // compiled and linked /ropi
```

See also

- *Creating Static Software Libraries with armar* for more information on library searching.

5.1.20 `__writeonly`

The `__writeonly` type qualifier indicates that a data object cannot be read from.

In the C and C++ type system it behaves as a cv-qualifier like **`const`** or **`volatile`**. Its specific effect is that an lvalue with `__writeonly` type cannot be converted to an rvalue.

Assignment to a `__writeonly` bitfield is not allowed if the assignment is implemented as read-modify-write. This is implementation-dependent.

Example

```
void foo(__writeonly int *ptr)
{
    *ptr = 0;                         // allowed
    printf("ptr value = %d\n", *ptr); // error
}
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.2 __declspec attributes

The __declspec keyword enables you to specify special attributes of objects and functions. For example, you can use the `__declspec` keyword to declare imported or exported functions and variables, or to declare *Thread Local Storage* (TLS) objects.

The `__declspec` keyword must prefix the declaration specification. For example:

```
__declspec(noreturn) void overflow(void);
__declspec(thread) int i;
```

Table 5-2 summarizes the available `__declspec` attributes. `__declspec` attributes are storage class modifiers. They do not affect the type of a function or variable.

Table 5-2 __declspec attributes supported by the compiler and their equivalents

__declspec attribute	non __declspec equivalent
`__declspec(dllexport)`	-
`__declspec(dllimport)`	-
`__declspec(noinline)`	`__attribute__((noinline))`[a]
`__declspec(noreturn)`	`__attribute__((noreturn))`[a]
`__declspec(nothrow)`	-
`__declspec(notshared)`	-
`__declspec(thread)`	-

a. A GNU compiler extension supported by the ARM compiler.

5.2.1 `__declspec(dllexport)`

The `__declspec(dllexport)` attribute exports the definition of a symbol through the dynamic symbol table when building DLL libraries. On classes, it controls the visibility of class impedimenta such as vtables, construction vtables and RTTI, and sets the default visibility for member function and static data members.

Usage

You can use `__declspec(dllexport)` on a function, a class, or on individual members of a class.

When an inline function is marked `__declspec(dllexport)`, the function definition might be inlined, but an out-of-line instance of the function is always generated and exported in the same way as for a non-inline function.

When a class is marked `__declspec(dllexport)`, for example,
`class __declspec(dllexport) S { ... };` its static data members and member functions are all exported. When individual static data members and member functions are marked with `__declspec(dllexport)`, only those members are exported. vtables, construction vtable tables and RTTI are also exported.

——— Note ———

The following declaration is correct:

```
class __declspec(dllexport) S { ... };
```

The following declaration is incorrect:

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

```
__declspec(dllexport) class S { ... };
```

In conjunction with `--export_all_vtbl`, you can use `__declspec(notshared)` to exempt a class or structure from having its vtable, construction vtable table and RTTI exported. `--export_all_vtbl` and `__declspec(dllexport)` are typically not used together.

Restrictions

If you mark a class with `__declspec(dllexport)`, you cannot then mark individual members of that class with `__declspec(dllexport)`.

If you mark a class with `__declspec(dllexport)`, ensure that all of the base classes of that class are marked `__declspec(dllexport)`.

If you export a virtual function within a class, ensure that you either export all of the virtual functions in that class, or that you define them inline so that they are visible to the client.

Example

The `__declspec()` required in a declaration depends on whether or not the definition is in the same shared library.

```
/* This is the declaration for use in the same shared library as the */
/* definition */
__declspec(dllexport) extern int mymod_get_version(void);

/* Translation unit containing the definition */
__declspec(dllexport) extern int mymod_get_version(void)
{
   return 42;
}

/* This is the declaration for use in a shared library that does not contain */
/* the definition */
__declspec(dllimport) extern int mymod_get_version(void);
```

As a result of the following macro, a translation unit that does not have the definition in a defining link unit sees `__declspec(dllexport)`.

```
/* mymod.h - interface to my module */
#ifdef BUILDING_MYMOD
#define MYMOD_API __declspec(dllexport)
#else /* not BUILDING_MYMOD */
#define MYMOD_API __declspec(dllimport)
#endif

MYMOD_API int mymod_get_version(void);
```

See also

- *__declspec(dllimport)* on page 5-21
- *__declspec(notshared)* on page 5-23
- *--export_all_vtbl, --no_export_all_vtbl* on page 3-38
- *--use_definition_visibility* on page 2-137 in the *Linker Reference*
- *--visibility_inlines_hidden* on page 3-95.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.2.2 `__declspec(dllimport)`

The `__declspec(dllimport)` attribute imports a symbol through the dynamic symbol table when linking against DLL libraries.

Usage

When an inline function is marked `__declspec(dllimport)`, the function definition in this compilation unit might be inlined, but is never generated out-of-line. An out-of-line call or address reference uses the imported symbol.

You can only use `__declspec(dllimport)` on **`extern`** functions and variables, and on classes.

When a class is marked `__declspec(dllimport)`, its static data members and member functions are all imported. When individual static data members and member functions are marked with `__declspec(dllimport)`, only those members are imported.

Restrictions

If you mark a class with `__declspec(dllimport)`, you cannot then mark individual members of that class with `__declspec(dllimport)`.

Examples

```
__declspec(dllimport) int i;
class __declspec(dllimport) X {   void f(); };
```

See also

- *__declspec(dllexport)* on page 5-19.

5.2.3 `__declspec(noinline)`

The `__declspec(noinline)` attribute suppresses the inlining of a function at the call points of the function.

`__declspec(noinline)` can also be applied to constant data, to prevent the compiler from using the value for optimization purposes, without affecting its placement in the object. This is a feature that can be used for patchable constants, that is, data that is later patched to a different value. It is an error to try to use such constants in a context where a constant value is required. For example, an array dimension.

Note

This `__declspec` attribute has the function attribute equivalent `__attribute__((noinline))`.

Examples

```
/* Prevent y being used for optimization */
__declspec(noinline) const int y = 5;
/* Suppress inlining of foo() wherever foo() is called */
__declspec(noinline) int foo(void);
```

See also

- *#pragma inline, #pragma no_inline* on page 5-54
- *__attribute__((noinline)) constant variable attribute* on page 5-42
- *__attribute__((noinline)) function attribute* on page 5-30.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.2.4 `__declspec(noreturn)`

The `__declspec(noreturn)` attribute asserts that a function never returns.

Note

This attribute has the function equivalent `__attribute__((noreturn))`. However, `__attribute__((noreturn))` and `__declspec(noreturn)` differ in that when compiling a function definition, if the function reaches an explicit or implicit return, `__attribute__((noreturn))` is ignored and the compiler generates a warning. This does not apply to `__declspec(noreturn)`.

Usage

Use this attribute to reduce the cost of calling a function that never returns, such as `exit()`. If a `noreturn` function returns to its caller, the behavior is undefined.

Restrictions

The return address is not preserved when calling the `noreturn` function. This limits the ability of a debugger to display the call stack.

Example

```
__declspec(noreturn) void overflow(void); // never return on overflow
int negate(int x)
{
    if (x == 0x80000000) overflow();
    return -x;
}
```

See also

- *__attribute__((noreturn)) function attribute* on page 5-31.

5.2.5 `__declspec(nothrow)`

The `__declspec(nothrow)` attribute asserts that a call to a function never results in a C++ exception being propagated from the call into the caller.

The ARM library headers automatically add this qualifier to declarations of C functions that, according to the ISO C Standard, can never throw.

Usage

If the compiler knows that a function can never throw out, it might be able to generate smaller exception-handling tables for callers of that function.

Restrictions

If a call to a function results in a C++ exception being propagated from the call into the caller, the behavior is undefined.

This modifier is ignored when not compiling with exceptions enabled.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example

```
struct S
{
    ~S();
};
__declspec(nothrow) extern void f(void);
void g(void)
{
    S s;
    f();
}
```

See also

- *--force_new_nothrow, --no_force_new_nothrow* on page 3-40
- *Using the ::operator new function* on page 6-11.

5.2.6 __declspec(notshared)

The `__declspec(notshared)` attribute prevents a specific class from having its virtual functions table and RTTI exported. This holds true regardless of other options you apply. For example, the use of `--export_all_vtbl` does not override `__declspec(notshared)`.

Example

```
struct __declspec(notshared) X
{
    virtual int f();
};                              // do not export this
int X::f()
{
    return 1;
}
struct Y : X
{
    virtual int g();
};                              // do export this
int Y::g()
{
    return 1;
}
```

5.2.7 __declspec(thread)

The `__declspec(thread)` attribute asserts that variables are thread-local and have *thread storage duration*, so that the linker arranges for the storage to be allocated automatically when a thread is created.

Note

The keyword `__thread` is supported as a synonym for `__declspec(thread)`.

Restrictions

File-scope thread-local variables cannot be dynamically initialized.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example

```
__declspec(thread) int i;
__thread int j;          // same as __decspec(thread) int j;
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.3 Function attributes

The `__attribute__` keyword enables you to specify special attributes of variables or structure fields, functions, and types. The keyword format is either:

```
__attribute__ ((attribute1, attribute2, ...))
__attribute__ ((__attribute1__, __attribute2__, ...))
```

For example:

```
void * Function_Attributes_malloc_0(int b) __attribute__ ((malloc));
static int b __attribute__ ((__unused__));
```

Table 5-3 summarizes the available function attributes.

Table 5-3 Function attributes supported by the compiler and their equivalents

Function attribute	non-attribute equivalent
`__attribute__((alias))`	-
`__attribute__((always_inline))`	`__forceinline`
`__attribute__((const))`	`__pure`
`__attribute__((constructor[(priority)]))`	-
`__attribute__((deprecated))`	-
`__attribute__((destructor[(priority)]))`	-
`__attribute__((format_arg(string-index)))`	-
`__attribute__((malloc))`	-
`__attribute__((noinline))`	`__declspec(noinline)`
`__attribute__((no_instrument_function))`	-
`__attribute__((nomerge))`	-
`__attribute__((nonnull))`	-
`__attribute__((noreturn))`	`__declspec(noreturn))`
`__attribute__((notailcall))`	-
`__attribute__((pcs("calling_convention")))`	-
`__attribute__((pure))`	-
`__attribute__((section("name")))`	-
`__attribute__((unused))`	-
`__attribute__((used))`	-
`__attribute__((visibility("visibility_type")))`	-
`__attribute__((weak))`	`__weak`
`__attribute__((weakref("target")))`	-

5.3.1 `__attribute__((alias))` function attribute

This function attribute enables you to specify multiple aliases for functions.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Where a function is defined in the current translation unit, the alias call is replaced by a call to the function, and the alias is emitted alongside the original name. Where a function is not defined in the current translation unit, the alias call is replaced by a call to the real function. Where a function is defined as **`static`**, the function name is replaced by the alias name and the function is declared external if the alias name is declared external.

Note

This function attribute is a GNU compiler extension supported by the ARM compiler.

Note

Variables names might also be aliased using the corresponding variable attribute `__attribute__((alias))`.

Syntax

return-type *newname*([argument-list]) __attribute__((alias("*oldname*")));

Where:

oldname is the name of the function to be aliased

newname is the new name of the aliased function.

Example

```
#include <stdio.h>
void foo(void)
{
    printf("%s\n", __FUNCTION__);
}
void bar(void) __attribute__((alias("foo")));
void gazonk(void)
{
    bar(); // calls foo
}
```

See also

- *__attribute__((alias)) variable attribute* on page 5-39.

5.3.2 `__attribute__((always_inline))` function attribute

This function attribute indicates that a function must be inlined.

The compiler attempts to inline the function, regardless of the characteristics of the function. However, the compiler does not inline a function if doing so causes problems. For example, a recursive function is inlined into itself only once.

Note

This function attribute is a GNU compiler extension that is supported by the ARM compiler. It has the keyword equivalent `__forceinline`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example

```
static int max(int x, int y) __attribute__((always_inline));
static int max(int x, int y)
{
    return x > y ? x : y; // always inline if possible
}
```

See also

- *--forceinline* on page 3-40
- *__forceinline* on page 5-5.

5.3.3 `__attribute__((const))` function attribute

Many functions examine only the arguments passed to them, and have no effects except for the return value. This is a much stricter class than `__attribute__((pure))`, because a function is not permitted to read global memory. If a function is known to operate only on its arguments then it can be subject to common sub-expression elimination and loop optimizations.

Note

This function attribute is a GNU compiler extension that is supported by the ARM compiler. It has the keyword equivalent `__pure`.

Example

```
int Function_Attributes_const_0(int b) __attribute__ ((const));
int Function_Attributes_const_0(int b)
{
    int aLocal=0;
    aLocal += Function_Attributes_const_0(b);
    aLocal += Function_Attributes_const_0(b);
    return aLocal;
}
```

In this code `Function_Attributes_const_0` might be called once only, with the result being doubled to obtain the correct return value.

See also

- *__attribute__((pure)) function attribute* on page 5-32
- *Functions that return the same result when called with the same arguments on page 5-24* in *Using the Compiler*.

5.3.4 `__attribute__((constructor[(priority)]))` function attribute

This attribute causes the function it is associated with to be called automatically before `main()` is entered.

Note

This attribute is a GNU compiler extension supported by the ARM compiler.

Syntax

```
__attribute__((constructor[(priority)]))
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Where *priority* is an optional integer value denoting the priority. A constructor with a low integer value runs before a constructor with a high integer value. A constructor with a priority runs before a constructor without a priority.

Priority values up to and including `100` are reserved for internal use. If you use these values, the compiler gives a warning. Priority values above `100` are not reserved.

Usage

You can use this attribute for start-up or initialization code. For example, to specify a function that is to be called when a DLL is loaded.

Example

In the following example, the constructor functions are called before execution enters `main()`, in the order specified:

```
int my_constructor(void) __attribute__((constructor));
int my_constructor2(void) __attribute__((constructor(102)));
int my_constructor3(void) __attribute__((constructor(101)));

int my_constructor(void) /* This is the 3rd constructor */
{                        /* function to be called */
    ...
    return 0;
}

int my_constructor2(void) /* This is the 1st constructor */
{                         /* function to be called */
    ...
    return 0;
}

int my_constructor3(void) /* This is the 2nd constructor */
{                         /* function to be called */
    ...
    return 0;
}
```

See also

- *__attribute__((destructor[(priority)])) function attribute* on page 5-29
- *--init=symbol* on page 2-63 in the *Linker Reference*

5.3.5 `__attribute__((deprecated))` function attribute

This function attribute indicates that a function exists but the compiler must generate a warning if the deprecated function is used.

Note

This function attribute is a GNU compiler extension that is supported by the ARM compiler.

Example

```
int Function_Attributes_deprecated_0(int b) __attribute__ ((deprecated));
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.3.6 `__attribute__((destructor[(priority)]))` function attribute

This attribute causes the function it is associated with to be called automatically after `main()` completes or after `exit()` is called.

Note

This attribute is a GNU compiler extension supported by the ARM compiler.

Syntax

```
__attribute__((destructor[(priority)]))
```

Where *`priority`* is an optional integer value denoting the priority. A destructor with a high integer value runs before a destructor with a low value. A destructor with a priority runs before a destructor without a priority.

Priority values up to and including `100` are reserved for internal use. If you use these values, the compiler gives a warning. Priority values above `100` are not reserved.

Example

```
int my_destructor(void) __attribute__((destructor));

int my_destructor(void) /* This function is called after main() */
{                       /* completes or after exit() is called. */
   ...
   return 0;
}
```

See also

- *__attribute__((constructor[(priority)])) function attribute* on page 5-27
- *--fini=symbol* on page 2-52 in the *Linker Reference*

5.3.7 `__attribute__((format_arg(string-index)))` function attribute

This function attribute specifies that a user-defined function modifies format strings. Use of this attribute enables calls to functions like `printf()`, `scanf()`, `strftime()` or `strfmon()`, whose operands are a call to the user-defined function, to be checked for errors.

Note

This function attribute is a GNU compiler extension that is supported by the ARM compiler.

5.3.8 `__attribute__((malloc))` function attribute

This function attribute indicates that the function can be treated like `malloc` and the compiler can perform the associated optimizations.

Note

This function attribute is a GNU compiler extension that is supported by the ARM compiler.

Example

```
void * Function_Attributes_malloc_0(int b) __attribute__ ((malloc));
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.3.9 `__attribute__((noinline))` function attribute

This function attribute suppresses the inlining of a function at the call points of the function.

——— **Note** ———

This function attribute is a GNU compiler extension that is supported by the ARM compiler. It has the `__declspec` equivalent `__declspec(noinline)`.

———————

Example

```
int fn(void) __attribute__((noinline));

int fn(void)
{
    return 42;
}
```

See also

- *#pragma inline, #pragma no_inline* on page 5-54
- *__attribute__((noinline)) constant variable attribute* on page 5-42
- *__declspec(noinline)* on page 5-21.

5.3.10 `__attribute__((nomerge))` function attribute

This function attribute prevents calls to the function that are distinct in the source from being combined in the object code.

See also

- *__attribute__((notailcall)) function attribute* on page 5-31
- *--retain=option* on page 3-83.

5.3.11 `__attribute__((nonnull))` function attribute

This function attribute specifies function parameters that are not supposed to be null pointers. This enables the compiler to generate a warning on encountering such a parameter.

——— **Note** ———

This function attribute is a GNU compiler extension that is supported by the ARM compiler.

———————

Syntax

```
__attribute__((nonnull(arg-index, ...)))
```

Where *`arg-index, ...`* denotes the argument index list.

If no argument index list is specified, all pointer arguments are marked as nonnull.

Example

The following declarations are equivalent:

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

```
void * my_memcpy (void *dest, const void *src, size_t len) __attribute__((nonnull (1,
2)));

void * my_memcpy (void *dest, const void *src, size_t len) __attribute__((nonnull));
```

5.3.12 `__attribute__((noreturn))` function attribute

This function attribute informs the compiler that the function does not return. The compiler can then perform optimizations by removing the code that is never reached.

Note

This function attribute is a GNU compiler extension that is supported by the ARM compiler. It has the `__declspec` equivalent `__declspec(noreturn)`. However, `__attribute((noreturn))` and `__declspec(noreturn)` differ in that when compiling a function definition, if the function reaches an explicit or implicit return, `__attribute((noreturn))` is ignored and the compiler generates a warning. This does not apply to `__declspec(noreturn)`.

Example

```
int Function_Attributes_NoReturn_0(void) __attribute__ ((noreturn));
```

See also

- *__declspec(noreturn)* on page 5-22.

5.3.13 `__attribute__((notailcall))` function attribute

This function attribute prevents tailcalling of the function. That is, the function is always called with a branch-and-link, even if (because the call occurs at the end of a function) the branch-and-link could be converted to a branch.

See also

- *__attribute__((nomerge)) function attribute* on page 5-30
- *--retain=option* on page 3-83.

5.3.14 `__attribute__((pcs("calling_convention")))`

This function attribute specifies the calling convention on targets with hardware floating-point, as an alternative to the **`__softfp`** keyword.

Note

This function attribute is a GNU compiler extension supported by the ARM compiler.

Syntax

`__attribute__((pcs("calling_convention")))`

Where *`calling_convention`* is one of the following:

`aapcs` uses integer registers, as for **`__softfp`**.

`aapcs-vfp` uses floating-point registers.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *__softfp* on page 5-12
- *Compiler support for floating-point computations and linkage on page 5-63* in *Using the Compiler*.

5.3.15 `__attribute__((pure))` function attribute

Many functions have no effects except to return a value, and their return value depends only on the parameters and global variables. Functions of this kind can be subject to data flow analysis and might be eliminated.

Note

This function attribute is a GNU compiler extension that is supported by the ARM compiler.

Although related, this function attribute is *not* equivalent to the `__pure` keyword. The function attribute equivalent to `__pure` is `__attribute__((const))`.

Example

```
int Function_Attributes_pure_0(int b) __attribute__ ((pure));
int Function_Attributes_pure_0(int b)
{
    return b++;

}

int foo(int b)
{
    int aLocal=0;
    aLocal += Function_Attributes_pure_0(b);
    aLocal += Function_Attributes_pure_0(b);
    return 0;
}
```

The call to `Function_Attributes_pure_0` in this example might be eliminated because its result is not used.

5.3.16 `__attribute__((section("name")))` function attribute

The `section` function attribute enables you to place code in different sections of the image.

Note

This function attribute is a GNU compiler extension that is supported by the ARM compiler.

Example

In the following example, `Function_Attributes_section_0` is placed into the RO section `new_section` rather than `.text`.

```
void Function_Attributes_section_0 (void)
    __attribute__ ((section ("new_section")));
void Function_Attributes_section_0 (void)
{
    static int aStatic =0;
    aStatic++;
}
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

In the following example, section function attribute overrides the #pragma arm section setting.

```
#pragma arm section code="foo"
  int f2()
  {
      return 1;
  }                                  // into the 'foo' area
  __attribute__ ((section ("bar"))) int f3()
  {
      return 1;
  }                                  // into the 'bar' area
  int f4()
  {
      return 1;
  }                                  // into the 'foo' area
#pragma arm section
```

See also

- *#pragma arm section [section_type_list]* on page 5-48.

5.3.17 `__attribute__((unused))` function attribute

The unused function attribute prevents the compiler from generating warnings if the function is not referenced. This does not change the behavior of the unused function removal process.

Note

This function attribute is a GNU compiler extension that is supported by the ARM compiler.

Example

```
static int Function_Attributes_unused_0(int b) __attribute__ ((unused));
```

5.3.18 `__attribute__((used))` function attribute

This function attribute informs the compiler that a static function is to be retained in the object file, even if it is unreferenced.

Static functions marked as used are emitted to a single section, in the order they are declared. You can specify the section functions are placed in using `__attribute__((section("name")))`.

Functions marked with `__attribute__((used))` are tagged in the object file to avoid removal by linker unused section removal.

Note

This function attribute is a GNU compiler extension that is supported by the ARM compiler.

Note

Static variables can also be marked as used using `__attribute__((used))`.

Example

```
static int lose_this(int);
static int keep_this(int) __attribute__((used));     // retained in object file
static int keep_this_too(int) __attribute__((used)); // retained in object file
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *__attribute__((section("name"))) function attribute* on page 5-32.
- *__attribute__((used)) variable attribute* on page 5-43
- *Elimination of unused sections* on page 5-4 in *Using the Linker*.

5.3.19 `__attribute__((visibility("visibility_type")))` function attribute

This function attribute affects the visibility of ELF symbols.

Note

This attribute is a GNU compiler extension supported by the ARM compiler.

Syntax

```
__attribute__((visibility("visibility_type")))
```

Where *visibility_type* is one of the following:

`default` The assumed visibility of symbols can be changed by other options. Default visibility overrides such changes. Default visibility corresponds to external linkage.

`hidden` The symbol is not placed into the dynamic symbol table, so no other executable or shared library can directly reference it. Indirect references are possible using function pointers.

`internal` Unless otherwise specified by the *processor-specific Application Binary Interface* (psABI), internal visibility means that the function is never called from another module.

`protected` The symbol is placed into the dynamic symbol table, but references within the defining module bind to the local symbol. That is, the symbol cannot be overridden by another module.

Usage

Except when specifying `default` visibility, this attribute is intended for use with declarations that would otherwise have external linkage.

You can apply this attribute to functions and variables in C and C++. In C++, it can also be applied to class, struct, union, and enum types, and namespace declarations.

Example

```
void __attribute__((visibility("internal"))) foo()
{
   ...
}
```

See also

- *--visibility_inlines_hidden* on page 3-95
- *--hide_all, --no_hide_all* on page 3-49
- *__attribute__((visibility("visibility_type"))) variable attribute* on page 5-44.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.3.20 `__attribute__((weak))` function attribute

Functions defined with `__attribute__((weak))` export their symbols weakly.

Functions declared with `__attribute__((weak))` and then defined without `__attribute__((weak))` behave as *weak* functions. This is not the same behavior as the `__weak` keyword.

Note

This function attribute is a GNU compiler extension that is supported by the ARM compiler.

Example

```
extern int Function_Attributes_weak_0 (int b) __attribute__ ((weak));
```

See also

- *__weak* on page 5-16.

5.3.21 `__attribute__((weakref("target")))` function attribute

This function attribute marks a function declaration as an alias that does not by itself require a function definition to be given for the target symbol.

Note

This function attribute is a GNU compiler extension supported by the ARM compiler.

Syntax

```
__attribute__((weakref("target")))
```

Where *`target`* is the target symbol.

Example

In the following example, `foo()` calls `y()` through a weak reference:

```
extern void y(void);
static void x(void) __attribute__((weakref("y")));
void foo (void)
{
  ...
  x();
  ...
}
```

Restrictions

This attribute can only be used on functions with internal linkage.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.4 Type attributes

The `__attribute__` keyword enables you to specify special attributes of variables or structure fields, functions, and types. The keyword format is either:

```
__attribute__ ((attribute1, attribute2, ...))
__attribute__ ((__attribute1__, __attribute2__, ...))
```

For example:

```
void * Function_Attributes_malloc_0(int b) __attribute__ ((malloc));
static int b __attribute__ ((__unused__));
```

Table 5-4 summarizes the available type attributes.

Table 5-4 Type attributes supported by the compiler and their equivalents

Type attribute	non-attribute equivalent
`__attribute__((bitband))`	-
`__attribute__((aligned))`	`__align`
`__attribute__((packed))`	`__packed`

5.4.1 `__attribute__((bitband))` type attribute

`__attribute__((bitband))` is a type attribute that gives you efficient atomic access to single-bit values in SRAM and Peripheral regions of the memory architecture. It is possible to set or clear a single bit directly with a single memory access in certain memory regions, rather than having to use the traditional read, modify, write approach. It is also possible to read a single bit directly rather than having to use the traditional read then shift and mask operation. Example 5-6 illustrates the use of `__attribute__((bitband))`.

Example 5-6 Using `__attribute__((bitband))`

```
typedef struct {
  int i : 1;
  int j : 2;
  int k : 3;
} BB __attribute__((bitband));

BB bb __attribute__((at(0x20000004));

void foo(void)
{
  bb.i = 1;
}
```

For peripherals that are sensitive to the memory access width, byte, halfword, and word stores or loads to the alias space are generated for **`char`**, **`short`**, and **`int`** types of bitfields of bit-banded structs respectively.

In Example 5-7 on page 5-37, bit-banded access is generated for `bb.i`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example 5-7 Bitfield bit-band access

```
typedef struct {
  char i : 1;
  int j : 2;
  int k : 3;
} BB __attribute__((bitband));

BB bb __attribute__((at(0x20000004)));

void foo()
{
  bb.i = 1;
}
```

If you do not use `__attribute__((at()))` to place the bit-banded variable in the bit-band region, you must relocate it using another method. You can do this by either using an appropriate scatter-loading description file or by using the `--rw_base` linker command-line option. See the *Linker Reference* for more information.

Restrictions

The following restrictions apply:

- This type attribute can only be used with **`struct`**. Any union type or other aggregate type with a union as a member cannot be bit-banded.
- Members of structs cannot be bit-banded individually.
- Bit-banded accesses are only generated for single-bit bitfields.
- Bit-banded accesses are not generated for **`const`** objects, pointers, and local objects.
- Bit-banding is only available on some processors. For example, the Cortex-M3 and Cortex-M4 processors.

See also

- *__attribute__((at(address))) variable attribute* on page 5-40
- *--bitband* on page 3-15
- the *Technical Reference Manual* for your processor.

5.4.2 `__attribute__((aligned))` type attribute

The `aligned` type attribute specifies a minimum alignment for the type.

Note

This type attribute is a GNU compiler extension that is supported by the ARM compiler.

5.4.3 `__attribute((packed))` type attribute

The `packed` type attribute specifies that a type must have the smallest possible alignment.

Note

This type attribute is a GNU compiler extension that is supported by the ARM compiler.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Errors

The compiler generates a warning message if you use this attribute in a `typedef`.

See also

- *__packed* on page 5-9
- *#pragma pack(n)* on page 5-56
- *Packed structures* on page 6-8
- *The __packed qualifier and unaligned data access in C and C++ code on page 5-46* in *Using the Compiler*
- *Detailed comparison of an unpacked struct, a __packed struct, and a struct with individually __packed fields on page 5-51* in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.5 Variable attributes

The `__attribute__` keyword enables you to specify special attributes of variables or structure fields, functions, and types. The keyword format is either:

```
__attribute__ ((attribute1, attribute2, ...))
__attribute__ ((__attribute1__, __attribute2__, ...))
```

For example:

```
void * Function_Attributes_malloc_0(int b) __attribute__ ((malloc));
static int b __attribute__ ((__unused__));
```

Table 5-3 on page 5-25 summarizes the available variable attributes.

Table 5-5 Variable attributes supported by the compiler and their equivalents

Variable attribute	non-attribute equivalent
`__attribute__((alias))`	-
`__attribute__((at(address)))`	-
`__attribute__((aligned))`	-
`__attribute__((deprecated))`	-
`__attribute__((noinline))`	
`__attribute__((packed))`	-
`__attribute__((section("name")))`	-
`__attribute__((unused))`	-
`__attribute__((used))`	-
`__attribute__((visibility("visibility_type")))`	-
`__attribute__((weak))`	`__weak`
`__attribute__((weakref("target")))`	
`__attribute__((zeroinit))`	-

5.5.1 `__attribute__((alias))` variable attribute

This variable attribute enables you to specify multiple aliases for variables.

Where a variable is defined in the current translation unit, the alias reference is replaced by a reference to the variable, and the alias is emitted alongside the original name. Where a variable is not defined in the current translation unit, the alias reference is replaced by a reference to the real variable. Where a variable is defined as **`static`**, the variable name is replaced by the alias name and the variable is declared external if the alias is declared external.

Note

Function names might also be aliased using the corresponding function attribute `__attribute__((alias))`.

Syntax

```
type newname __attribute__((alias("oldname")));
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Where:

oldname is the name of the variable to be aliased

newname is the new name of the aliased variable.

Example

```
#include <stdio.h>
int oldname = 1;
extern int newname __attribute__((alias("oldname"))); // declaration
void foo(void)
{
    printf("newname = %d\n", newname); // prints 1
}
```

See also

- *__attribute__((alias)) function attribute* on page 5-25.

5.5.2 `__attribute__((at(address)))` variable attribute

This variable attribute enables you to specify the absolute address of a variable.

The variable is placed in its own section, and the section containing the variable is given an appropriate type by the compiler:

- Read-only variables are placed in a section of type RO.
- Initialized read-write variables are placed in a section of type RW.

  Variables explicitly initialized to zero are placed in:
  - A section of type ZI in RVCT 4.0 and later.
  - A section of type RW (not ZI) in RVCT 3.1 and earlier. Such variables are not candidates for the ZI-to-RW optimization of the compiler.
- Uninitialized variables are placed in a section of type ZI.

Note

This variable attribute is not supported by GNU compilers.

Syntax

`__attribute__((at(address)))`

Where:

address is the desired address of the variable.

Restrictions

The linker is not always able to place sections produced by the `at` variable attribute.

The compiler faults use of the `at` attribute when it is used on declarations with incomplete types.

Errors

The linker gives an error message if it is not possible to place a section at a specified address.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Examples

```
const int x1 __attribute__((at(0x10000))) = 10; /* RO */
int x2 __attribute__((at(0x12000))) = 10; /* RW */
int x3 __attribute__((at(0x14000))) = 0; /* RVCT 3.1 and earlier: RW.
                                          * RVCT 4.0 and later: ZI. */
int x4 __attribute__((at(0x16000))); /* ZI */
```

See also

- *Using __at sections to place sections at a specific address* on page 8-35 in *Using the Linker*.

5.5.3 `__attribute__((aligned))` variable attribute

The `aligned` variable attribute specifies a minimum alignment for the variable or structure field, measured in bytes.

Note

This variable attribute is a GNU compiler extension that is supported by the ARM compiler.

Examples

```
/* Aligns on 16-byte boundary */
int x __attribute__ ((aligned (16)));

/* In this case, the alignment used is the maximum alignment for a scalar data type.
For ARM, this is 8 bytes. */
short my_array[3] __attribute__ ((aligned));
```

See also

- *__align* on page 5-2.

5.5.4 `__attribute__((deprecated))` variable attribute

The `deprecated` variable attribute enables the declaration of a deprecated variable without any warnings or errors being issued by the compiler. However, any access to a `deprecated` variable creates a warning but still compiles. The warning gives the location where the variable is used and the location where it is defined. This helps you to determine why a particular definition is deprecated.

Note

This variable attribute is a GNU compiler extension that is supported by the ARM compiler.

Example

```
extern int Variable_Attributes_deprecated_0 __attribute__ ((deprecated));
extern int Variable_Attributes_deprecated_1 __attribute__ ((deprecated));
void Variable_Attributes_deprecated_2()
{
    Variable_Attributes_deprecated_0=1;
    Variable_Attributes_deprecated_1=2;
}
```

Compiling this example generates two warning messages.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.5.5 `__attribute__((noinline))` constant variable attribute

The `noinline` variable attribute prevents the compiler from making any use of a constant data value for optimization purposes, without affecting its placement in the object. This feature can be used for patchable constants, that is, data that is later patched to a different value. It is an error to try to use such constants in a context where a constant value is required. For example, an array dimension.

Example

```
__attribute__((noinline)) const int m = 1;
```

See also

- *#pragma inline, #pragma no_inline* on page 5-54
- *__attribute__((noinline)) function attribute* on page 5-30
- *__declspec(noinline)* on page 5-21.

5.5.6 `__attribute__((packed))` variable attribute

The `packed` variable attribute specifies that a variable or structure field has the smallest possible alignment. That is, one byte for a variable, and one bit for a field, unless you specify a larger value with the `aligned` attribute.

Note

This variable attribute is a GNU compiler extension that is supported by the ARM compiler.

Example

```
struct
{
    char a;
    int b __attribute__ ((packed));
} Variable_Attributes_packed_0;
```

See also

- *__packed* on page 5-9
- *#pragma pack(n)* on page 5-56
- *Packed structures* on page 6-8
- *The __packed qualifier and unaligned data access in C and C++ code on page 5-46* in *Using the Compiler*
- *Detailed comparison of an unpacked struct, a __packed struct, and a struct with individually __packed fields on page 5-51* in *Using the Compiler*.

5.5.7 `__attribute__((section("name")))` variable attribute

Normally, the ARM compiler places the objects it generates in sections like `.data` and `.bss`. However, you might require additional data sections or you might want a variable to appear in a special section, for example, to map to special hardware. The `section` attribute specifies that a variable must be placed in a particular data section. If you use the `section` attribute, read-only variables are placed in RO data sections, read-write variables are placed in RW data sections unless you use the `zero_init` attribute. In this case, the variable is placed in a ZI section.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Note

This variable attribute is a GNU compiler extension supported by the ARM compiler.

Example

```
/* in RO section */
const int descriptor[3] __attribute__ ((section ("descr"))) = { 1,2,3 };

/* in RW section */
long long rw_initialized[10] __attribute__ ((section ("INITIALIZED_RW"))) = {5};

/* in RW section */
long long rw[10] __attribute__ ((section ("RW")));

/* in ZI section */
long long altstack[10] __attribute__ ((section ("STACK"), zero_init));
```

See also

- *How to find where a symbol is placed when linking* on page 6-6 in *Using the Linker*
- *Using fromelf to find where a symbol is placed in an executable ELF image* on page 3-7.

5.5.8 `__attribute__((unused))` variable attribute

Normally, the compiler warns if a variable is declared but is never referenced. This attribute informs the compiler that you expect a variable to be unused and tells it not issue a warning if it is not used.

Note

This variable attribute is a GNU compiler extension that is supported by the ARM compiler.

Example

```
void Variable_Attributes_unused_0()
{
    static int aStatic =0;
    int aUnused __attribute__ ((unused));
    int bUnused;
    aStatic++;
}
```

In this example, the compiler warns that `bUnused` is declared but never referenced, but does not warn about `aUnused`.

Note

The GNU compiler does not give any warning.

5.5.9 `__attribute__((used))` variable attribute

This variable attribute informs the compiler that a static variable is to be retained in the object file, even if it is unreferenced.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Static variables marked as used are emitted to a single section, in the order they are declared. You can specify the section that variables are placed in using `__attribute__((section("name")))`.

Data marked with `__attribute__((used))` is tagged in the object file to avoid removal by linker unused section removal.

Note

This variable attribute is a GNU compiler extension that is supported by the ARM compiler.

Note

Static functions can also be marked as used using `__attribute__((used))`.

Usage

You can use `__attribute__((used))` to build tables in the object.

Example

```
static int lose_this = 1;
static int keep_this __attribute__((used)) = 2;     // retained in object file
static int keep_this_too __attribute__((used)) = 3; // retained in object file
```

See also

- *__attribute__((section("name"))) variable attribute* on page 5-42
- *__attribute__((used)) function attribute* on page 5-33
- *Elimination of unused sections* on page 5-4 in *Using the Linker*.

5.5.10 `__attribute__((visibility("visibility_type")))` variable attribute

This variable attribute affects the visibility of ELF symbols.

Note

This attribute is a GNU compiler extension supported by the ARM compiler.

Syntax

`__attribute__((visibility("visibility_type")))`

Where *`visibility_type`* is one of the following:

`default` The assumed visibility of symbols can be changed by other options. Default visibility overrides such changes. Default visibility corresponds to external linkage.

`hidden` The symbol is not placed into the dynamic symbol table, so no other executable or shared library can directly reference it. Indirect references are possible using function pointers.

`internal` Unless otherwise specified by the *processor-specific Application Binary Interface* (psABI), internal visibility means that the function is never called from another module.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

`protected` The symbol is placed into the dynamic symbol table, but references within the defining module bind to the local symbol. That is, the symbol cannot be overridden by another module.

Usage

Except when specifying `default` visibility, this attribute is intended for use with declarations that would otherwise have external linkage.

You can apply this attribute to functions and variables in C and C++. In C++, you can also apply it to class, struct, union, and enum types, and namespace declarations.

Example

```
int i __attribute__((visibility("hidden")));
```

See also

- *--hide_all, --no_hide_all* on page 3-49
- *__attribute__((visibility("visibility_type"))) function attribute* on page 5-34.

5.5.11 `__attribute__((weak))` variable attribute

The declaration of a weak variable is permitted, and acts in a similar way to `__weak`. The equivalent is:

```
__weak int Variable_Attributes_weak_compare;
```

Note

This variable attribute is a GNU compiler extension that is supported by the ARM compiler.

See also

- *__weak* on page 5-16.

5.5.12 `__attribute__((weakref("target")))` variable attribute

This variable attribute marks a variable declaration as an alias that does not by itself require a definition to be given for the target symbol.

Note

This variable attribute is a GNU compiler extension supported by the ARM compiler.

Syntax

```
__attribute__((weakref("target")))
```

Where *`target`* is the target symbol.

Example

In the following example, `a` is assigned the value of `y` through a weak reference:

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

```
extern int y;
static int x __attribute__((weakref("y")));

void foo (void)
{
  int a = x;
  ...
}
```

Restrictions

This attribute can only be used on variables that are declared as static.

5.5.13 `__attribute__((zero_init))` variable attribute

The `section` attribute specifies that a variable must be placed in a particular data section. The `zero_init` attribute specifies that a variable with no initializer is placed in a ZI data section. If an initializer is specified, an error is reported.

Example

```
__attribute__((zero_init)) int x;                    /* in section ".bss" */
__attribute__((section("mybss"), zero_init)) int y;  /* in section "mybss" */
```

See also

- *__attribute__((section("name"))) variable attribute* on page 5-42.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.6 Pragmas

The ARM compiler recognizes a number of ARM-specific pragmas. Table 5-6 summarizes the available pragmas.

Note

Pragmas override related command-line options. For example, `#pragma arm` overrides the command-line option `--thumb`.

Table 5-6 Pragmas supported by the compiler

Pragmas		
#pragma anon_unions, #pragma no_anon_unions	#pragma hdrstop	#pragma pack(*n*)
#pragma arm	#pragma import *symbol_name*	#pragma pop
#pragma arm section [*section_type_list*]	#pragma import(__use_full_stdio)	#pragma push
#pragma diag_default *tag*[,*tag*,...]	#pragma import(__use_smaller_memcpy)	#pragma softfp_linkage, no_softfp_linkage
#pragma diag_error *tag*[,*tag*,...]	#pragma inline, #pragma no_inline	#pragma unroll [(*n*)]
#pragma diag_remark *tag*[,*tag*,...]	#pragma no_pch	#pragma unroll_completely
#pragma diag_suppress *tag*[,*tag*,...]	#pragma O*num*	#pragma thumb
#pragma diag_warning *tag*[,*tag*,...]	#pragma once	#pragma weak *symbol*
#pragma [no_]exceptions_unwind	#pragma Ospace	#pragma weak *symbol1* = *symbol2*
#pragma GCC system_header	#pragma Otime	

5.6.1 #pragma anon_unions, #pragma no_anon_unions

These pragmas enable and disable support for anonymous structures and unions.

Default

The default is `#pragma no_anon_unions`.

See also

- *Anonymous classes, structures and unions* on page 4-16

5.6.2 #pragma arm

This pragma switches code generation to the ARM instruction set. It overrides the `--thumb` compiler option.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--arm* on page 3-11
- *--thumb* on page 3-90
- *#pragma thumb* on page 5-60.

5.6.3 `#pragma arm section [section_type_list]`

This pragma specifies a section name to be used for subsequent functions or objects. This includes definitions of anonymous objects the compiler creates for initializations.

Note

You can use `__attribute__((section(..)))` for functions or variables as an alternative to `#pragma arm section`.

Syntax

`#pragma arm section [section_type_list]`

Where:

section_type_list specifies an optional list of section names to be used for subsequent functions or objects. The syntax of *section_type_list* is:

`section_type[[=]"name"] [,section_type="name"]*`

Valid section types are:

- `code`
- `rodata`
- `rwdata`
- `zidata`.

Usage

Use `#pragma arm section [section_type_list]` to place functions and variables in separate named sections. The scatter-loading description file can then be used to locate these at a particular address in memory.

Restrictions

This option has no effect on:

- Inline functions and their local static variables.
- Template instantiations and their local static variables.
- Elimination of unused variables and functions. However, using `#pragma arm section` might enable the linker to eliminate a function or variable that might otherwise be kept because it is in the same section as a used function or variable.
- The order that definitions are written to the object file.

Example

```
int x1 = 5;                     // in .data (default)
int y1[100];                    // in .bss (default)
int const z1[3] = {1,2,3};      // in .constdata (default)
#pragma arm section rwdata = "foo", rodata = "bar"
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

```
int x2 = 5;                      // in foo (data part of region)
int y2[100];                     // in .bss
int const z2[3] = {1,2,3};       // in bar
char *s2 = "abc";                // s2 in foo, "abc" in .conststring
#pragma arm section rodata
int x3 = 5;                      // in foo
int y3[100];                     // in .bss
int const z3[3] = {1,2,3};       // in .constdata
char *s3 = "abc";                // s3 in foo, "abc" in .conststring
#pragma arm section code = "foo"
int add1(int x)                    // in foo (code part of region)
{
    return x+1;
}
#pragma arm section code
```

See also

- *__attribute__((section("name"))) function attribute* on page 5-32
- Chapter 8 *Using scatter files* in *Using the Linker*.

5.6.4 `#pragma diag_default` *tag[,tag,...]*

This pragma returns the severity of the diagnostic messages that have the specified tags to the severities that were in effect before any pragmas were issued. Diagnostic messages are messages whose message numbers are postfixed by `-D`, for example, `#550-D`.

Syntax

`#pragma diag_default` *tag*`[,`*tag*`,...]`

Where:

tag`[,`*tag*`,...]` is a comma-separated list of diagnostic message numbers specifying the messages whose severities are to be changed.

At least one diagnostic message number must be specified.

Example

```
// <stdio.h> not #included deliberately
#pragma diag_error 223
void hello(void)
{
    printf("Hello ");
}
#pragma diag_default 223
void world(void)
{
    printf("world!\n");
}
```

Compiling this code with the option `--diag_warning=223` generates diagnostic messages to report that the function `printf()` is declared implicitly.

The effect of `#pragma diag_default 223` is to return the severity of diagnostic message 223 to Warning severity, as specified by the `--diag_warning` command-line option.

See also

- *--diag_warning=tag[,tag,...]* on page 3-33

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *#pragma diag_error tag[,tag,...]*
- *#pragma diag_remark tag[,tag,...]*
- *#pragma diag_suppress tag[,tag,...]* on page 5-51
- *#pragma diag_warning tag[, tag, ...]* on page 5-51
- *Compiler diagnostics* on page 6-2 in *Using the Compiler*.

5.6.5 `#pragma diag_error` *tag[,tag,...]*

This pragma sets the diagnostic messages that have the specified tags to Error severity. Diagnostic messages are messages whose message numbers are postfixed by `-D`, for example, `#550-D`.

Syntax

`#pragma diag_error` *tag*`[,`*tag*`,...]`

Where:

`tag[,tag,...]` is a comma-separated list of diagnostic message numbers specifying the messages whose severities are to be changed.

At least one diagnostic message number must be specified.

See also

- *--diag_error=tag[,tag,...]* on page 3-30
- *#pragma diag_default tag[,tag,...]* on page 5-49
- *#pragma diag_remark tag[,tag,...]*
- *#pragma diag_suppress tag[,tag,...]* on page 5-51
- *#pragma diag_warning tag[, tag, ...]* on page 5-51
- *Options that change the severity of compiler diagnostic messages* on page 6-4 in *Using the Compiler*.

5.6.6 `#pragma diag_remark` *tag[,tag,...]*

This pragma sets the diagnostic messages that have the specified tags to Remark severity. Diagnostic messages are messages whose message numbers are postfixed by `-D`, for example, `#550-D`.

`#pragma diag_remark` behaves analogously to `#pragma diag_errors`, except that the compiler sets the diagnostic messages having the specified tags to Remark severity rather than Error severity.

Note

Remarks are not displayed by default. Use the `--remarks` compiler option to see remark messages.

Syntax

`#pragma diag_remark` *tag*`[,`*tag*`,...]`

Where:

`tag[,tag,...]` is a comma-separated list of diagnostic message numbers specifying the messages whose severities are to be changed.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--diag_remark=tag[,tag,...]* on page 3-31
- *--remarks* on page 3-82
- *#pragma diag_default tag[,tag,...]* on page 5-49
- *#pragma diag_error tag[,tag,...]* on page 5-50
- *#pragma diag_suppress tag[,tag,...]*
- *#pragma diag_warning tag[, tag, ...]*
- *Options that change the severity of compiler diagnostic messages* on page 6-4 in *Using the Compiler*.

5.6.7 `#pragma diag_suppress tag[,tag,...]`

This pragma disables all diagnostic messages that have the specified tags. Diagnostic messages are messages whose message numbers are postfixed by `-D`, for example, `#550-D`.

`#pragma diag_suppress` behaves analogously to `#pragma diag_errors`, except that the compiler suppresses the diagnostic messages having the specified tags rather than setting them to have Error severity.

Syntax

```
#pragma diag_suppress tag[,tag,...]
```

Where:

`tag[,tag,...]` is a comma-separated list of diagnostic message numbers specifying the messages to be suppressed.

See also

- *--diag_suppress=tag[,tag,...]* on page 3-32
- *#pragma diag_default tag[,tag,...]* on page 5-49
- *#pragma diag_error tag[,tag,...]* on page 5-50
- *#pragma diag_remark tag[,tag,...]* on page 5-50
- *#pragma diag_warning tag[, tag, ...]*
- Chapter 6 *Compiler Diagnostic Messages* in *Using the Compiler*.

5.6.8 `#pragma diag_warning tag[, tag, ...]`

This pragma sets the diagnostic messages that have the specified tags to Warning severity. Diagnostic messages are messages whose message numbers are postfixed by `-D`, for example, `#550-D`.

`#pragma diag_warning` behaves analogously to `#pragma diag_errors`, except that the compiler sets the diagnostic messages having the specified tags to Warning severity rather than Error severity.

Syntax

```
#pragma diag_warning tag[,tag,...]
```

Where:

`tag[,tag,...]` is a comma-separated list of diagnostic message numbers specifying the messages whose severities are to be changed.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *--diag_warning=tag[,tag,...]* on page 3-33
- *#pragma diag_default tag[,tag,...]* on page 5-49
- *#pragma diag_error tag[,tag,...]* on page 5-50
- *#pragma diag_remark tag[,tag,...]* on page 5-50
- *#pragma diag_suppress tag[,tag,...]* on page 5-51
- *Options that change the severity of compiler diagnostic messages* on page 6-4 in *Using the Compiler*.

5.6.9 `#pragma exceptions_unwind`, `#pragma no_exceptions_unwind`

These pragmas enable and disable function unwinding at runtime.

Default

The default is `#pragma exceptions_unwind`.

See also

- *--exceptions, --no_exceptions* on page 3-37
- *--exceptions_unwind, --no_exceptions_unwind* on page 3-37
- *Function unwinding at runtime* on page 6-15.

5.6.10 `#pragma hdrstop`

This pragma enables you to specify where the set of precompilation header files end.

This pragma must appear before the first token that does not belong to a preprocessing directive.

See also

- *PreCompiled Header (PCH) files* on page 4-29 in *Using the Compiler*.

5.6.11 `#pragma import` *`symbol_name`*

This pragma generates an importing reference to *`symbol_name`*. This is the same as the assembler directive:

```
IMPORT symbol_name
```

Syntax

```
#pragma import symbol_name
```

Where:

`symbol_name` is a symbol to be imported.

Usage

You can use this pragma to select certain features of the C library, such as the heap implementation or real-time division. If a feature described in this book requires a symbol reference to be imported, the required symbol is specified.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *Using the C library with an application* on page 2-33 in *Using ARM® C and C++ Libraries and Floating-Point Support*.

5.6.12 `#pragma import(__use_full_stdio)`

This pragma selects an extended version of microlib that uses full standard ANSI C input and output functionality.

Note

Microlib is an alternative library to the default C library. Only use this pragma if you are using microlib.

The following exceptions apply:

- `feof()` and `ferror()` always return `0`
- `setvbuf()` and `setbuf()` are guaranteed to fail.

`feof()` and `ferror()` always return `0` because the error and end-of-file indicators are not supported.

`setvbuf()` and `setbuf()` are guaranteed to fail because all streams are unbuffered.

This version of microlib `stdio` can be retargeted in the same way as the standardlib `stdio` functions.

See also

- *--library_type=lib* on page 3-58
- *About microlib* on page 3-2 in *Using ARM® C and C++ Libraries and Floating-Point Support*
- *Tailoring input/output functions in the C and C++ libraries* on page 2-92 in *Using ARM® C and C++ Libraries and Floating-Point Support*.

5.6.13 `#pragma import(__use_smaller_memcpy)`

This pragma selects a smaller, but slower, version of `memcpy()` for use with the C micro-library (microlib). A byte-by-byte implementation of `memcpy()` using `LDRB` and `STRB` is used.

Note

Microlib is an alternative library to the default C library. Only use this pragma if you are using microlib.

Default

The default version of `memcpy()` used by microlib is a larger, but faster, word-by-word implementation using `LDR` and `STR`.

See also

- *--library_type=lib* on page 3-58
- Chapter 3 *The ARM C micro-library* in *Using ARM® C and C++ Libraries and Floating-Point Support*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.6.14 `#pragma inline`, `#pragma no_inline`

These pragmas control inlining, similar to the `--inline` and `--no_inline` command-line options. A function defined under `#pragma no_inline` is not inlined into other functions, and does not have its own calls inlined.

The effect of suppressing inlining into other functions can also be achieved by marking the function as `__declspec(noinline)` or `__attribute__((noinline))`.

Default

The default is `#pragma inline`.

See also

- *--inline, --no_inline* on page 3-53
- *__declspec(noinline)* on page 5-21
- *__attribute__((noinline)) constant variable attribute* on page 5-42
- *__attribute__((noinline)) function attribute* on page 5-30.

5.6.15 `#pragma no_pch`

This pragma suppresses PCH processing for a given source file.

See also

- *--pch* on page 3-75
- *PreCompiled Header (PCH) files* on page 4-29 in *Using the Compiler*.

5.6.16 `#pragma O`*`num`*

This pragma changes the optimization level.

Syntax

`#pragma O`*`num`*

Where:

`num` is the new optimization level.

The value of *`num`* is 0, 1, 2 or 3.

Usage

This pragma enables you to assign optimization levels to individual functions.

Restriction

The pragma must be placed outside the function.

See also

- *-Onum* on page 3-71
- *#pragma Ospace* on page 5-55
- *#pragma Otime* on page 5-55.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.6.17 `#pragma once`

This pragma enables the compiler to skips subsequent includes of that header file.

`#pragma once` is accepted for compatibility with other compilers, and enables you to use other forms of header guard coding. However, it is preferable to use `#ifndef` and `#define` coding because this is more portable.

Example

The following example shows the placement of a `#ifndef` guard around the body of the file, with a `#define` of the guard variable after the `#ifndef`.

```
#ifndef FILE_H
#define FILE_H
#pragma once        // optional ... body of the header file ...#endif
```

The `#pragma once` is marked as optional in this example. This is because the compiler recognizes the `#ifndef` header guard coding and skips subsequent includes even if `#pragma once` is absent.

5.6.18 `#pragma Ospace`

This pragma instructs the compiler to perform optimizations to reduce image size at the expense of a possible increase in execution time.

Usage

This pragma enables you to assign optimization goals to individual functions.

Restriction

The pragma must be placed outside the function.

See also

- *-Ospace* on page 3-72
- *#pragma Onum* on page 5-54
- *#pragma Otime*.

5.6.19 `#pragma Otime`

This pragma instructs the compiler to perform optimizations to reduce execution time at the expense of a possible increase in image size.

Usage

This pragma enables you to assign optimization goals to individual functions.

Restriction

The pragma must be placed outside the function.

See also

- *-Otime* on page 3-73
- *#pragma Onum* on page 5-54
- *#pragma Ospace*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.6.20 `#pragma pack(n)`

This pragma aligns members of a structure to the minimum of `n` and their natural alignment. Packed objects are read and written using unaligned accesses.

Syntax

`#pragma pack(n)`

Where:

n is the alignment in bytes, valid alignment values being `1`, `2`, `4` and `8`.

Default

The default is `#pragma pack(8)`.

Example

This example demonstrates how `pack(2)` aligns integer variable `b` to a 2-byte boundary.

```
typedef struct
{
    char a;
    int b;
} S;

#pragma pack(2)

typedef struct
{
    char a;
    int b;
} SP;

S var = { 0x11, 0x44444444 };
SP pvar = { 0x11, 0x44444444 };
```

The layout of `S` is as shown in Figure 5-1, while the layout of `SP` is as shown in Figure 5-2. In Figure 5-2, `x` denotes one byte of padding.

0 a	1	2 padding	3
4 b	5 b	6 b	7 b

Figure 5-1 Nonpacked structure S

0 a	1 x	2 b	3 b
4 b	5 b		

Figure 5-2 Packed structure SP

Note

`SP` is a 6-byte structure. There is no padding after `b`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *__packed* on page 5-9
- *__attribute__((packed)) variable attribute* on page 5-42
- *Packed structures* on page 6-8
- *The __packed qualifier and unaligned data access in C and C++ code on page 5-46* in *Using the Compiler*
- *Detailed comparison of an unpacked struct, a __packed struct, and a struct with individually __packed fields on page 5-51* in *Using the Compiler*.

5.6.21 `#pragma pop`

This pragma restores the previously saved pragma state.

See also

- *#pragma push*.

5.6.22 `#pragma push`

This pragma saves the current pragma state.

See also

- *#pragma pop*.

5.6.23 `#pragma softfp_linkage`, `#pragma no_softfp_linkage`

These pragmas control software floating-point linkage.

`#pragma softfp_linkage` asserts that all function declarations up to the next `#pragma no_softfp_linkage` describe functions that use software floating-point linkage.

Note

This pragma has the keyword equivalent `__softfp`.

Usage

This pragma can be useful when applied to an entire interface specification, located in the header file, without altering that file.

Default

The default is `#pragma no_softfp_linkage`.

See also

- *__softfp* on page 5-12
- *Compiler support for floating-point computations and linkage on page 5-63* in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.6.24 `#pragma unroll [(n)]`

This pragma instructs the compiler to unroll a loop by *n* iterations.

Syntax

`#pragma unroll`

`#pragma unroll (n)`

Where:

n is an optional value indicating the number of iterations to unroll.

Default

If you do not specify a value for *n*, the compiler assumes `#pragma unroll (4)`.

Usage

This pragma is only applicable if you are compiling with `-O3 -Otime`. When compiling with `-O3 -Otime`, the compiler automatically unrolls loops where it is beneficial to do so. You can use this pragma to ask the compiler to unroll a loop that has not been unrolled automatically.

Note

Use this pragma only when you have evidence, for example from `--diag_warning=optimizations`, that the compiler is not unrolling loops optimally by itself.

You cannot determine whether this pragma is having any effect unless you compile with `--diag_warning=optimizations` or examine the generated assembly code, or both.

Restrictions

This pragma can only take effect when you compile with `-O3 -Otime`. Even then, the use of this pragma is a *request* to the compiler to unroll a loop that has not been unrolled automatically. It does not guarantee that the loop is unrolled.

`#pragma unroll [(n)]` can be used only immediately before a **`for`** loop, a **`while`** loop, or a **`do`** ... **`while`** loop.

Example

```
void matrix_multiply(float ** __restrict dest, float ** __restrict src1,
    float ** __restrict src2, unsigned int n)
{
    unsigned int i, j, k;
    for (i = 0; i < n; i++)
    {
        for (k = 0; k < n; k++)
        {
            float sum = 0.0f;
            /* #pragma unroll */
            for(j = 0; j < n; j++)
                sum += src1[i][j] * src2[j][k];
            dest[i][k] = sum;
        }
    }
}
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

In this example, the compiler does not normally complete its loop analysis because `src2` is indexed as `src2[j][k]` but the loops are nested in the opposite order, that is, with `j` inside `k`. When `#pragma unroll` is uncommented in the example, the compiler proceeds to unroll the loop four times.

If the intention is to multiply a matrix that is not a multiple of four in size, for example an *n* * *n* matrix, `#pragma unroll (m)` might be used instead, where *m* is some value so that *n* is an integral multiple of *m*.

See also

- *--diag_warning=optimizations* on page 3-34
- *-Onum* on page 3-71
- *-Otime* on page 3-73
- *#pragma unroll_completely*
- *Loop unrolling in C code on page 5-11* in *Using the Compiler*.

5.6.25 `#pragma unroll_completely`

This pragma instructs the compiler to completely unroll a loop. It has an effect only if the compiler can determine the number of iterations the loop has.

Usage

This pragma is only applicable if you are compiling with `-O3 -Otime`. When compiling with `-O3 -Otime`, the compiler automatically unrolls loops where it is beneficial to do so. You can use this pragma to ask the compiler to completely unroll a loop that has not automatically been unrolled completely.

Note

Use this `#pragma` only when you have evidence, for example from `--diag_warning=optimizations`, that the compiler is not unrolling loops optimally by itself.

You cannot determine whether this pragma is having any effect unless you compile with `--diag_warning=optimizations` or examine the generated assembly code, or both.

Restrictions

This pragma can only take effect when you compile with `-O3 -Otime`. Even then, the use of this pragma is a *request* to the compiler to unroll a loop that has not been unrolled automatically. It does not guarantee that the loop is unrolled.

`#pragma unroll_completely` can only be used immediately before a **`for`** loop, a **`while`** loop, or a **`do`** ... **`while`** loop.

Using `#pragma unroll_completely` on an outer loop can prevent vectorization. On the other hand, using `#pragma unroll_completely` on an inner loop might help in some cases.

See also

- *--diag_warning=optimizations* on page 3-34
- *-Onum* on page 3-71
- *-Otime* on page 3-73
- *#pragma unroll [(n)]* on page 5-58
- *Loop unrolling in C code on page 5-11* in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.6.26 `#pragma thumb`

This pragma switches code generation to the Thumb instruction set. It overrides the `--arm` compiler option.

If you are compiling code for a pre-Thumb-2 processor and using VFP, *any* function containing floating-point operations is compiled for ARM.

See also

- *--arm* on page 3-11
- *--thumb* on page 3-90
- *#pragma arm* on page 5-47.

5.6.27 `#pragma weak symbol`, `#pragma weak symbol1 = symbol2`

This pragma is a deprecated language extension to mark symbols as weak or to define weak aliases of symbols. It is an alternative to using the **`__weak`** keyword or the GCC `weak` and `alias` attributes.

Example

In the following example, `weak_fn` is declared as a weak alias of `__weak_fn`:

```
extern void weak_fn(int a);
#pragma weak weak_fn = __weak_fn

void __weak_fn(int a)
{
    ...
}
```

See also

- *__attribute__((alias)) variable attribute* on page 5-39
- *__attribute__((weak)) function attribute* on page 5-35
- *__attribute__((weak)) variable attribute* on page 5-45
- *__weak* on page 5-16.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.7 Instruction intrinsics

This section describes instruction intrinsics for realizing ARM machine language instructions from C or C++ code. Table 5-7 summarizes the available intrinsics.

Table 5-7 Instruction intrinsics supported by the ARM compiler

Instruction intrinsics		
`__breakpoint`	`__ldrt`	`__schedule_barrier`
`__cdp`	`__memory_changed`	`__semihost`
`__clrex`	`__nop`	`__sev`
`__clz`	`__pld`	`__sqrt`
`__current_pc`	`__pldw`	`__sqrtf`
`__current_sp`	`__pli`	`__ssat`
`__disable_fiq`	`__promise`	`__strex`
`__disable_irq`	`__qadd`	`__strexd`
`__enable_fiq`	`__qdbl`	`__strt`
`__enable_irq`	`__qsub`	`__swp`
`__fabs`	`__rbit`	`__usat`
`__fabsf`	`__rev`	`__wfe`
`__force_stores`	`__return_address`	`__wfi`
`__ldrex`	`__ror`	`__yield`
`__ldrexd`		

5.7.1 `__breakpoint` intrinsic

This intrinsic inserts a `BKPT` instruction into the instruction stream generated by the compiler. It enables you to include a breakpoint instruction in your C or C++ code.

Syntax

```
void __breakpoint(int val)
```

Where:

val is a compile-time constant integer whose range is:

`0 ... 65535`	if you are compiling source as ARM code
`0 ... 255`	if you are compiling source as Thumb code.

Errors

The compiler does not recognize the `__breakpoint` intrinsic when compiling for a target that does not support the `BKPT` instruction. The compiler generates either a warning or an error in this case.

The undefined instruction trap is taken if a `BKPT` instruction is executed on an architecture that does not support it.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example

```
void func(void)
{
    ...
    __breakpoint(0xF02C);
    ...
}
```

See also

- *BKPT* on page 3-134 in the *Assembler Reference*.

5.7.2 __cdp intrinsic

This intrinsic inserts a `CDP` or `CDP2` instruction into the instruction stream generated by the compiler. It enables you to include coprocessor data operations in your C or C++ code.

——— **Note** ———

This intrinsic is intended for specialist expert use only.

———————————

Syntax

`__cdp(unsigned int *coproc*, unsigned int *ops*, unsigned int *regs*)`

Where:

coproc Identifies the coprocessor the instruction is for.

coproc must be an integer in the range 0 to 15.

ops Is an encoding of two opcodes where:

- the first opcode is a 4-bit coprocessor-specific opcode
- the second opcode is an optional 3-bit coprocessor-specific opcode.

Add `0x100` to *ops* to generate a `CDP2` instruction.

regs Is an encoding of the coprocessor registers.

Usage

The use of these instructions depends on the coprocessor. See your coprocessor documentation for more information.

Example

```
void copro_example()
{
    const unsigned int ops = 0xA3; // opcode1 = A, opcode2 = 3
    const unsigned int regs = 0xCDE; // reg1 = C, reg2 = D, reg3 = E
    __cdp(4,ops,regs); // coprocessor number 4
}
```

See also

- *CDP and CDP2* on page 3-125 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.7.3 `__clrex` intrinsic

This intrinsic inserts a `CLREX` instruction into the instruction stream generated by the compiler. It enables you to include a `CLREX` instruction in your C or C++ code.

Syntax

```
void __clrex(void)
```

Errors

The compiler does not recognize the `__clrex` intrinsic when compiling for a target that does not support the `CLREX` instruction. The compiler generates either a warning or an error in this case.

See also

- *CLREX* on page 3-42 in the *Assembler Reference*.

5.7.4 `__clz` intrinsic

This intrinsic inserts a `CLZ` instruction or an equivalent code sequence into the instruction stream generated by the compiler. It enables you to count the number of leading zeros of a data value in your C or C++ code.

Syntax

```
unsigned char __clz(unsigned int val)
```

Where:

`val` is an **`unsigned int`**.

Return value

The `__clz` intrinsic returns the number of leading zeros in *`val`*.

See also

- *CLZ* on page 3-58 in the *Assembler Reference*.

5.7.5 `__current_pc` intrinsic

This intrinsic enables you to determine the current value of the program counter at the point in your program where the intrinsic is used.

Syntax

```
unsigned int __current_pc(void)
```

Return value

The `__current_pc` intrinsic returns the current value of the program counter at the point in the program where the intrinsic is used.

See also

- *__current_sp intrinsic* on page 5-64
- *__return_address intrinsic* on page 5-77

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *Legacy inline assembler that accesses sp, lr, or pc* on page 7-56 in *Using the Compiler*.

5.7.6 `__current_sp` intrinsic

This intrinsic returns the value of the stack pointer at the current point in your program.

Syntax

```
unsigned int __current_sp(void)
```

Return value

The `__current_sp` intrinsic returns the current value of the stack pointer at the point in the program where the intrinsic is used.

See also

- *__current_pc intrinsic* on page 5-63
- *__return_address intrinsic* on page 5-77
- *Legacy inline assembler that accesses sp, lr, or pc* on page 7-56 in *Using the Compiler*.

5.7.7 `__disable_fiq` intrinsic

This intrinsic disables FIQ interrupts.

Note

Typically, this intrinsic disables FIQ interrupts by setting the F-bit in the CPSR. However, for v7-M it sets the fault mask register (FAULTMASK). FIQ interrupts are not supported in v6-M.

Syntax

```
int __disable_fiq(void);
```

```
void __disable_fiq(void);
```

Usage

`int __disable_fiq(void);` disables fast interrupts and returns the value the FIQ interrupt mask has in the PSR prior to the disabling of interrupts.

`void __disable_fiq(void);` disables fast interrupts.

Return value

`int __disable_fiq(void);` returns the value the FIQ interrupt mask has in the `PSR` prior to the disabling of FIQ interrupts.

Restrictions

`int __disable_fiq(void);` is not supported when compiling with `--cpu=7`. This is because of the difference between the generic ARMv7 architecture and the ARMv7 A, R, and M-profiles in the exception handling model. This means that when you compile with `--cpu=7`, the compiler is unable to generate an instruction sequence that works on all ARMv7 processors and therefore `int __disable_fiq(void);` is not supported. You can use the `void __disable_fiq(void);` function prototype with `--cpu=7`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

The `__disable_fiq` intrinsic can only be executed in privileged modes, that is, in non-user modes. In User mode this intrinsic does not change the interrupt flags in the CPSR.

Example

```
void foo(void)
{
    int was_masked = __disable_fiq();
    /* ... */
    if (!was_masked)
        __enable_fiq();
}
```

See also

- *__enable_fiq intrinsic* on page 5-66.

5.7.8 `__disable_irq` intrinsic

This intrinsic disables IRQ interrupts.

Note

Typically, this intrinsic disables IRQ interrupts by setting the I-bit in the CPSR. However, for M-profile it sets the exception mask register (PRIMASK).

Syntax

`int __disable_irq(void);`

`void __disable_irq(void);`

Usage

`int __disable_irq(void);` disables interrupts and returns the value the IRQ interrupt mask has in the PSR prior to the disabling of interrupts.

`void __disable_irq(void);` disables interrupts.

Return value

`int __disable_irq(void);` returns the value the IRQ interrupt mask has in the PSR prior to the disabling of IRQ interrupts.

Example

```
void foo(void)
{
    int was_masked = __disable_irq();
    /* ... */
    if (!was_masked)
        __enable_irq();
}
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Restrictions

`int __disable_irq(void);` is not supported when compiling with `--cpu=7`. This is because of the difference between the generic ARMv7 architecture and the ARMv7 A, R, and M-profiles in the exception handling model. This means that when you compile with `--cpu=7`, the compiler is unable to generate an instruction sequence that works on all ARMv7 processors and therefore `int __disable_irq(void);` is not supported. You can use the `void __disable_irq(void);` function prototype with `--cpu=7`.

The following example illustrates the difference between compiling for ARMv7-M and ARMv7-R:

```
/* test.c */ void DisableIrq(void)
{
  __disable_irq();
} int DisableIrq2(void)
{
  return __disable_irq();
}

armcc -c --cpu=Cortex-M3 -o m3.o test.c

  DisableIrq
    0x00000000:  b672      r.    CPSID    i
    0x00000002:  4770      pG    BX       lr
  DisableIrq2
    0x00000004:  f3ef8010  ....  MRS      r0,PRIMASK
    0x00000008:  f0000001  ....  AND      r0,r0,#1
    0x0000000c:  b672      r.    CPSID    i
    0x0000000e:  4770      pG    BX       lr


armcc -c --cpu=Cortex-R4 --thumb -o r4.o test.c

  DisableIrq
    0x00000000:  b672      r.    CPSID    i
    0x00000002:  4770      pG    BX       lr
  DisableIrq2
    0x00000004:  f3ef8000  ....  MRS      r0,APSR ; formerly CPSR
    0x00000008:  f00000080 ....  AND      r0,r0,#0x80
    0x0000000c:  b672      r.    CPSID    i
    0x0000000e:  4770      pG    BX       lr
```

In all cases, the `__disable_irq` intrinsic can only be executed in privileged modes, that is, in non-user modes. In User mode this intrinsic does not change the interrupt flags in the `CPSR`.

See also

- *__enable_irq intrinsic* on page 5-67.

5.7.9 `__enable_fiq` intrinsic

This intrinsic enables FIQ interrupts.

Note

Typically, this intrinsic enables FIQ interrupts by clearing the F-bit in the CPSR. However, for v7-M, it clears the fault mask register (FAULTMASK). FIQ interrupts are not supported in v6-M.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Syntax

```
void __enable_fiq(void)
```

Restrictions

The `__enable_fiq` intrinsic can only be executed in privileged modes, that is, in non-user modes. In User mode this intrinsic does not change the interrupt flags in the `CPSR`.

See also

- *__disable_fiq intrinsic* on page 5-64.

5.7.10 `__enable_irq` intrinsic

This intrinsic enables IRQ interrupts.

Note

Typically, this intrinsic enables IRQ interrupts by clearing the I-bit in the CPSR. However, for Cortex M-profile processors, it clears the exception mask register (PRIMASK).

Syntax

```
void __enable_irq(void)
```

Restrictions

The `__enable_irq` intrinsic can only be executed in privileged modes, that is, in non-user modes. In User mode this intrinsic does not change the interrupt flags in the `CPSR`.

See also

- *__disable_irq intrinsic* on page 5-65.

5.7.11 `__fabs` intrinsic

This intrinsic inserts a `VABS` instruction or an equivalent code sequence into the instruction stream generated by the compiler. It enables you to obtain the absolute value of a double-precision floating-point value from within your C or C++ code.

Note

The `__fabs` intrinsic is an analogue of the standard C library function `fabs`. It differs from the standard library function in that a call to `__fabs` is guaranteed to be compiled into a single, inline, machine instruction on an ARM architecture-based processor equipped with a VFP coprocessor.

Syntax

```
double __fabs(double val)
```

Where:

`val` is a double-precision floating-point value.

Return value

The `__fabs` intrinsic returns the absolute value of *`val`* as a **`double`**.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *__fabsf intrinsic*
- e in the *Assembler Reference*.

5.7.12 `__fabsf` intrinsic

This intrinsic is a single-precision version of the `__fabs` intrinsic. It is functionally equivalent to `__fabs`, except that:

- it takes an argument of type **`float`** instead of an argument of type **`double`**
- it returns a **`float`** value instead of a **`double`** value.

Syntax

```
float __fabs(float val)
```

See also

- *__fabs intrinsic* on page 5-67

5.7.13 `__force_stores` intrinsic

This intrinsic causes all variables that are visible outside the current function, such as variables that have pointers to them passed into or out of the function, to be written back to memory if they have been changed.

This intrinsic also acts as a scheduling barrier.

Syntax

```
void __force_stores(void)
```

See also

- *__memory_changed intrinsic* on page 5-72
- *__schedule_barrier intrinsic* on page 5-78.

5.7.14 `__ldrex` intrinsic

This intrinsic inserts an instruction of the form `LDREX[size]` into the instruction stream generated by the compiler. It enables you to load data from memory in your C or C++ code using an `LDREX` instruction. *`size`* in `LDREX[size]` is `B` for byte stores or `H` for halfword stores. If no size is specified, word stores are performed.

Syntax

```
unsigned int __ldrex(volatile void *ptr)
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Where:

ptr points to the address of the data to be loaded from memory. To specify the type of the data to be loaded, cast the parameter to an appropriate pointer type.

Table 5-8 Access widths supported by the __ldrex intrinsic

Instruction	Size of data loaded	C cast
LDREXB	unsigned byte	(**unsigned char** *)
LDREXB	signed byte	(**signed char** *)
LDREXH	unsigned halfword	(**unsigned short** *)
LDREXH	signed halfword	(**short** *)
LDREX	word	(**int** *)

Return value

The `__ldrex` intrinsic returns the data loaded from the memory address pointed to by *ptr*.

Errors

The compiler does not recognize the `__ldrex` intrinsic when compiling for a target that does not support the `LDREX` instruction. The compiler generates either a warning or an error in this case.

The `__ldrex` intrinsic does not support access to doubleword data. The compiler generates an error if you specify an access width that is not supported.

Example

```
int foo(void)
{
    int loc = 0xff;
    return __ldrex((volatile char *)loc);
}
```

Compiling this code with the command-line option `--cpu=6k` produces

```
||foo|| PROC
    MOV      r0,#0xff
    LDREXB   r0,[r0]
    BX       lr
    ENDP
```

See also

- *__ldrexd intrinsic*
- *__strex intrinsic* on page 5-81
- *__strexd intrinsic* on page 5-82
- *LDREX and STREX* on page 3-39 in the *Assembler Reference*.

5.7.15 `__ldrexd` intrinsic

This intrinsic inserts an `LDREXD` instruction into the instruction stream generated by the compiler. It enables you to load data from memory in your C or C++ code using an `LDREXD` instruction. It supports access to doubleword data.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Syntax

```
unsigned long long __ldrexd(volatile void *ptr)
```

Where:

`ptr` points to the address of the data to be loaded from memory. To specify the type of the data to be loaded, cast the parameter to an appropriate pointer type.

Table 5-9 Access widths supported by the __ldrex intrinsic

Instruction	Size of data loaded	C cast
`LDREXD`	unsigned long long	`(unsigned long long *)`
`LDREXD`	signed long long	`(signed long long *)`

Return value

The `__ldrexd` intrinsic returns the data loaded from the memory address pointed to by *ptr*.

Errors

The compiler does not recognize the `__ldrexd` intrinsic when compiling for a target that does not support the `LDREXD` instruction. The compiler generates either a warning or an error in this case.

The `__ldrexd` intrinsic only supports access to doubleword data. The compiler generates an error if you specify an access width that is not supported.

See also

- *__ldrex intrinsic* on page 5-68
- *__strex intrinsic* on page 5-81
- *__strexd intrinsic* on page 5-82
- *LDREX and STREX* on page 3-39 in the *Assembler Reference*.

5.7.16 `__ldrt` intrinsic

This intrinsic inserts an assembly language instruction of the form `LDR{size}T` into the instruction stream generated by the compiler. It enables you to load data from memory in your C or C++ code using an `LDRT` instruction.

Syntax

```
unsigned int __ldrt(const volatile void *ptr)
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Where:

`ptr` Points to the address of the data to be loaded from memory. To specify the size of the data to be loaded, cast the parameter to an appropriate integral type.

Table 5-10 Access widths supported by the __ldrt intrinsic

Instruction[a]	Size of data loaded	C cast
LDRSBT	signed byte	(**signed char** *)
LDRBT	unsigned byte	(**char** *)
LDRSHT	signed halfword	(**signed short int** *)
LDRHT	unsigned halfword	(**short int** *)
LDRT	word	(**int** *)

a. Or equivalent.

Return value

The `__ldrt` intrinsic returns the data loaded from the memory address pointed to by *ptr*.

Errors

The compiler does not recognize the `__ldrt` intrinsic when compiling for a target that does not support the `LDRT` instruction. The compiler generates either a warning or an error in this case.

The `__ldrt` intrinsic does not support access to doubleword data. The compiler generates an error if you specify an access width that is not supported.

Example

```
int foo(void)
{
    int loc = 0xff;
    return __ldrt((const volatile int *)loc);
}
```

Compiling this code with the default options produces:

```
||foo|| PROC
    MOV      r0,#0xff
    LDRBT    r1,[r0],#0
    MOV      r2,#0x100
    LDRBT    r0,[r2],#0
    ORR      r0,r1,r0,LSL #8
    BX       lr
    ENDP
```

See also

- *--thumb* on page 3-90
- *LDR and STR, unprivileged* on page 3-17 in the *ARM Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.7.17 `__memory_changed` intrinsic

This intrinsic causes all variables that are visible outside the current function, such as variables that have pointers to them passed into or out of the function, to be written back to memory if they have been changed, and then to be read back from memory.

This intrinsic also acts as a scheduling barrier.

Syntax

```
void __memory_changed(void)
```

See also

- *__force_stores intrinsic* on page 5-68
- *__schedule_barrier intrinsic* on page 5-78.

5.7.18 `__nop`

This intrinsic inserts a `NOP` instruction or an equivalent code sequence into the instruction stream generated by the compiler. One `NOP` instruction is generated for each `__nop` intrinsic in the source.

The compiler does not optimize away the `NOP` instructions, except for normal unreachable code elimination. The `__nop` intrinsic also acts as a barrier for instruction scheduling in the compiler. That is, instructions are not moved from one side of the `NOP` to the other as a result of optimization.

Note

You can use the `__schedule_barrier` intrinsic to insert a scheduling barrier without generating a `NOP` instruction.

Syntax

```
void __nop(void)
```

See also

- *__sev intrinsic* on page 5-79
- *__schedule_barrier intrinsic* on page 5-78
- *__wfe intrinsic* on page 5-86
- *__wfi intrinsic* on page 5-86
- *__yield intrinsic* on page 5-87
- *NOP* on page 3-143 in the *Assembler Reference*
- *Generic intrinsics supported by the compiler* on page 4-7 in *Using the Compiler*.

5.7.19 `__pld` intrinsic

This intrinsic inserts a data prefetch, for example `PLD`, into the instruction stream generated by the compiler. It enables you to signal to the memory system from your C or C++ program that a data load from an address is likely in the near future.

Syntax

```
void __pld(...)
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Where:

`...` denotes any number of pointer or integer arguments specifying addresses of memory to prefetch.

Restrictions

If the target architecture does not support data prefetching, the compiler generates neither a `PLD` instruction nor a `NOP` instruction, but ignores the intrinsic.

Example

```
extern int data1;
extern int data2;
volatile int *interrupt = (volatile int *)0x8000;
volatile int *uart = (volatile int *)0x9000;
void get(void)
{
    __pld(data1, data2);
    while (!*interrupt);
    *uart = data1;        // trigger uart as soon as interrupt occurs
    *(uart+1) = data2;
}
```

See also

- *__pldw intrinsic*
- *__pli intrinsic* on page 5-74
- *PLD, PLDW, and PLI* on page 3-28 in the *Assembler Reference*.

5.7.20 `__pldw` intrinsic

This intrinsic inserts a `PLDW` instruction into the instruction stream generated by the compiler. It enables you to signal to the memory system from your C or C++ program that a data load from an address with an intention to write is likely in the near future.

Syntax

```
void __pldw(...)
```

Where:

`...` denotes any number of pointer or integer arguments specifying addresses of memory to prefetch.

Restrictions

If the target architecture does not support data prefetching, this intrinsic has no effect.

This intrinsic only takes effect in ARMv7 architectures and above that provide Multiprocessing Extensions. That is, when the predefined macro `__TARGET_FEATURE_MULTIPROCESSING` is defined.

Example

```
void foo(int *bar)
{
     __pldw(bar);
}
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

See also

- *Compiler predefines* on page 5-98
- *__pld intrinsic* on page 5-72
- *__pli intrinsic*
- *PLD, PLDW, and PLI* on page 3-28 in the *Assembler Reference*.

5.7.21 `__pli` intrinsic

This intrinsic inserts an instruction prefetch, for example `PLI`, into the instruction stream generated by the compiler. It enables you to signal to the memory system from your C or C++ program that an instruction load from an address is likely in the near future.

Syntax

```
void __pli(...)
```

Where:

`...` denotes any number of pointer or integer arguments specifying addresses of instructions to prefetch.

Restrictions

If the target architecture does not support instruction prefetching, the compiler generates neither a `PLI` instruction nor a `NOP` instruction, but ignores the intrinsic.

See also

- *__pld intrinsic* on page 5-72
- *__pldw intrinsic* on page 5-73
- *PLD, PLDW, and PLI* on page 3-28 in the *Assembler Reference*.

5.7.22 `__promise` intrinsic

This intrinsic represents a promise you make to the compiler that a given expression always has a nonzero value. This enables the compiler to perform more aggressive optimization when vectorizing code.

Syntax

```
void __promise(expr)
```

Where *expr* is an expression that evaluates to nonzero.

Usage

`__promise(expr)` is similar to `assert()`. However, unlike `assert()`, `__promise(expr)` is effective even when `NDEBUG` is defined.

If assertions are enabled (by including `assert.h` and not defining `NDEBUG`) then the promise is additionally checked at runtime using `assert()`.

expr is not to have side-effects, and is not evaluated unless `assert.h` is included and `NDEBUG` is not defined.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.7.23 `__qadd` intrinsic

This intrinsic inserts a `QADD` instruction into the instruction stream generated by the compiler. It enables you to obtain the saturating add of two integers from within your C or C++ code.

Note

The compiler might optimize your code when it detects opportunity to do so, using equivalent instructions from the same family to produce fewer instructions.

Syntax

```
int __qadd(int val1, int val2)
```

Where:

val1 is the first summand of the saturating add operation

val2 is the second summand of the saturating add operation.

Return value

The `__qadd` intrinsic returns the saturating add of *val1* and *val2*.

Restriction

This intrinsic is only available on targets that have the `QADD` instruction.

See also

- *__qdbl intrinsic*
- *__qsub intrinsic* on page 5-76
- *QADD, QSUB, QDADD, and QDSUB* on page 3-97 in the *Assembler Reference*.

5.7.24 `__qdbl` intrinsic

This intrinsic inserts instructions equivalent to the saturating add of an integer with itself into the instruction stream generated by the compiler. It enables you to obtain the saturating double of an integer from within your C or C++ code.

Syntax

```
int __qdbl(int val)
```

Where:

val is the data value to be doubled.

Return value

The `__qdbl` intrinsic returns the saturating add of *val* with itself, or equivalently, `__qadd(val, val)`.

See also

- *__qadd intrinsic*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.7.25 `__qsub` intrinsic

This intrinsic inserts a `QSUB` instruction or an equivalent code sequence into the instruction stream generated by the compiler. It enables you to obtain the saturating subtraction of two integers from within your C or C++ code.

Syntax

```
int __qsub(int val1, int val2)
```

Where:

val1 is the minuend of the saturating subtraction operation

val2 is the subtrahend of the saturating subtraction operation.

Return value

The `__qsub` intrinsic returns the saturating subtraction of *val1* and *val2*.

See also

- *__qadd intrinsic* on page 5-75
- *QADD, QSUB, QDADD, and QDSUB* on page 3-97 in the *Assembler Reference*.

5.7.26 `__rbit` intrinsic

This intrinsic inserts an `RBIT` instruction into the instruction stream generated by the compiler. It enables you to reverse the bit order in a 32-bit word from within your C or C++ code.

Syntax

```
unsigned int __rbit(unsigned int val)
```

where:

val is the data value whose bit order is to be reversed.

Return value

The `__rbit` intrinsic returns the value obtained from `val` by reversing its bit order.

See also

- *REV, REV16, REVSH, and RBIT* on page 3-69 in the *Assembler Reference*.

5.7.27 `__rev` intrinsic

This intrinsic inserts a `REV` instruction or an equivalent code sequence into the instruction stream generated by the compiler. It enables you to convert a 32-bit big-endian data value into a little-endian data value, or a 32-bit little-endian data value into a big-endian data value from within your C or C++ code.

Note

The `__rev` intrinsic is available irrespective of the target processor or architecture you are compiling for. However, if the `REV` instruction is not available on the target, the compiler compensates with an alternative code sequence that could increase the number of instructions, effectively expanding the intrinsic into a function.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Note

The compiler introduces `REV` automatically when it recognizes certain expressions.

Syntax

```
unsigned int __rev(unsigned int val)
```

Where:

val is an **unsigned int**.

Return value

The `__rev` intrinsic returns the value obtained from `val` by reversing its byte order.

See also

- *REV, REV16, REVSH, and RBIT* on page 3-69 in the *Assembler Reference*.

5.7.28 `__return_address` intrinsic

This intrinsic returns the return address of the current function.

Syntax

```
unsigned int __return_address(void)
```

Return value

The `__return_address` intrinsic returns the value of the link register that is used in returning from the current function.

Restrictions

The `__return_address` intrinsic does *not* affect the ability of the compiler to perform optimizations such as inlining, tailcalling, and code sharing. Where optimizations are made, the value returned by `__return_address` reflects the optimizations performed:

No optimization

When no optimizations are performed, the value returned by `__return_address` from within a function `foo` is the return address of `foo`.

Inline optimization

If a function `foo` is inlined into a function `bar` then the value returned by `__return_address` from within `foo` is the return address of `bar`.

Tail-call optimization

If a function `foo` is tail-called from a function `bar` then the value returned by `__return_address` from within `foo` is the return address of `bar`.

See also

- *__current_pc intrinsic* on page 5-63
- *__current_sp intrinsic* on page 5-64
- *Legacy inline assembler that accesses sp, lr, or pc* on page 7-56 in the *Compiler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.7.29 `__ror` intrinsic

This intrinsic inserts a `ROR` instruction or operand rotation into the instruction stream generated by the compiler. It enables you to rotate a value right by a specified number of places from within your C or C++ code.

Note

The compiler introduces `ROR` automatically when it recognizes certain expressions.

Syntax

```
unsigned int __ror(unsigned int val, unsigned int shift)
```

Where:

`val` is the value to be shifted right

`shift` is a constant shift in the range 1-31.

Return value

The `__ror` intrinsic returns the value of `val` rotated right by `shift` number of places.

See also

- *ASR, LSL, LSR, ROR, and RRX* on page 3-71 in the *Assembler Reference*.

5.7.30 `__schedule_barrier` intrinsic

This intrinsic creates a sequence point where operations before and operations after the sequence point are not merged by the compiler. A scheduling barrier does not cause memory to be updated. If variables are held in registers they are updated in place, and not written out.

This intrinsic is similar to the `__nop` intrinsic, except that no `NOP` instruction is generated.

Syntax

```
void __schedule_barrier(void)
```

See also

- *__nop* on page 5-72

5.7.31 `__semihost` intrinsic

This intrinsic inserts an `SVC` or `BKPT` instruction into the instruction stream generated by the compiler. It enables you to make semihosting calls from C or C++ that are independent of the target architecture.

Syntax

```
int __semihost(int val, const void *ptr)
```

Where:

`val` Is the request code for the semihosting request.

`ptr` Is a pointer to an argument/result block.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Return value

See *Developing Software for ARM® Processors* for more information on the results of semihosting calls.

Usage

Use this intrinsic from C or C++ to generate the appropriate semihosting call for your target and instruction set:

`SVC 0x123456`	In ARM state for all architectures.
`SVC 0xAB`	In Thumb state, excluding ARMv7-M. This behavior is not guaranteed on *all* debug targets from ARM or from third parties.
`BKPT 0xAB`	For ARMv7-M, Thumb-2 only.

Restrictions

ARM processors prior to ARMv7 use `SVC` instructions to make semihosting calls. However, if you are compiling for a Cortex M-profile processor, semihosting is implemented using the `BKPT` instruction.

Example

```
char buffer[100];
...
void foo(void)
{
    __semihost(0x01, (const void *)buf); // equivalent in thumb state to
                                          // int __svc(0xAB) my_svc(int, int *);
                                          // result = my_svc(0x1, &buffer);
}
```

Compiling this code with the option `--thumb` generates:

```
||foo|| PROC
    ...
    LDR      r1,|L1.12|
    MOVS     r0,#1
    SVC      #0xab
    ...
|L1.12|
    ...
buffer
    %        400
```

See also

- *--cpu=list* on page 3-20
- *--thumb* on page 3-90
- *__svc* on page 5-13
- *BKPT* on page 3-134 in the *Assembler Reference*
- *SVC* on page 3-135 in the *Assembler Reference*

5.7.32 `__sev` intrinsic

This intrinsic inserts a `SEV` instruction into the instruction stream generated by the compiler.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

In some architectures, for example the v6T2 architecture, the `SEV` instruction executes as a `NOP` instruction.

Syntax

```
void __sev(void)
```

Errors

The compiler does not recognize the `__sev` intrinsic when compiling for a target that does not support the `SEV` instruction. The compiler generates either a warning or an error in this case.

See also

- *__nop* on page 5-72
- *__wfe intrinsic* on page 5-86
- *__wfi intrinsic* on page 5-86
- *__yield intrinsic* on page 5-87
- *NOP* on page 3-143 in the *Assembler Reference*.

5.7.33 `__sqrt` intrinsic

This intrinsic inserts a VFP `VSQRT` instruction into the instruction stream generated by the compiler. It enables you to obtain the square root of a double-precision floating-point value from within your C or C++ code.

Note

The `__sqrt` intrinsic is an analogue of the standard C library function `sqrt`. It differs from the standard library function in that a call to `__sqrt` is guaranteed to be compiled into a single, inline, machine instruction on an ARM architecture-based processor equipped with a VFP coprocessor.

Syntax

```
double __sqrt(double val)
```

Where:

`val` is a double-precision floating-point value.

Return value

The `__sqrt` intrinsic returns the square root of *`val`* as a **`double`**.

Errors

The compiler does not recognize the `__sqrt` intrinsic when compiling for a target that is not equipped with a VFP coprocessor. The compiler generates either a warning or an error in this case.

See also

- *__sqrtf intrinsic* on page 5-81
- *VABS, VNEG, and VSQRT* on page 4-8 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.7.34 __sqrtf intrinsic

This intrinsic is a single-precision version of the __sqrtf intrinsic. It is functionally equivalent to __sqrt, except that:

- it takes an argument of type **float** instead of an argument of type **double**
- it returns a **float** value instead of a **double** value.

See also

- *__sqrt intrinsic* on page 5-80
- *VABS, VNEG, and VSQRT* on page 4-8 in the *Assembler Reference*.

5.7.35 __ssat intrinsic

This intrinsic inserts an SSAT instruction into the instruction stream generated by the compiler. It enables you to saturate a signed value from within your C or C++ code.

Syntax

```
int __ssat(int val, unsigned int sat)
```

Where:

val Is the value to be saturated.

sat Is the bit position to saturate to.

sat must be in the range 1 to 32.

Return value

The __ssat intrinsic returns *val* saturated to the signed range $-2^{sat-1} \leq x \leq 2^{sat-1} - 1$.

Errors

The compiler does not recognize the __ssat intrinsic when compiling for a target that does not support the SSAT instruction. The compiler generates either a warning or an error in this case.

See also

- *__usat intrinsic* on page 5-85
- *SSAT and USAT* on page 3-99 in the *Assembler Reference*.

5.7.36 __strex intrinsic

This intrinsic inserts an instruction of the form STREX[size] into the instruction stream generated by the compiler. It enables you to use an STREX instruction in your C or C++ code to store data to memory.

Syntax

```
int __strex(unsigned int val, volatile void *ptr)
```

Where:

val is the value to be written to memory.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

ptr points to the address of the data to be written to in memory. To specify the size of the data to be written, cast the parameter to an appropriate integral type.

Table 5-11 Access widths supported by the `__strex` intrinsic

Instruction	Size of data stored	C cast
STREXB	unsigned byte	(**char** *)
STREXH	unsigned halfword	(**short** *)
STREX	word	(**int** *)

Return value

The `__strex` intrinsic returns:

`0` if the STREX instruction succeeds

`1` if the STREX instruction is locked out.

Errors

The compiler does not recognize the `__strex` intrinsic when compiling for a target that does not support the STREX instruction. The compiler generates either a warning or an error in this case.

The `__strex` intrinsic does not support access to doubleword data. The compiler generates an error if you specify an access width that is not supported.

Example

```
int foo(void)
{
    int loc=0xff;
    return(!__strex(0x20, (volatile char *)loc));
}
```

Compiling this code with the command-line option `--cpu=6k` produces

```
||foo|| PROC
    MOV      r0,#0xff
    MOV      r2,#0x20
    STREXB   r1,r2,[r0]
    RSBS     r0,r1,#1
    MOVCC    r0,#0
    BX       lr
ENDP
```

See also

- *__ldrex intrinsic* on page 5-68
- *__ldrexd intrinsic* on page 5-69
- *__strexd intrinsic*
- *LDREX and STREX* on page 3-39 in the *Assembler Reference*.

5.7.37 `__strexd` intrinsic

This intrinsic inserts an STREXD instruction into the instruction stream generated by the compiler. It enables you to use an STREXD instruction in your C or C++ code to store data to memory. It supports exclusive stores of doubleword data to memory.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Syntax

```
int __strexd(unsigned long long val, volatile void *ptr)
```

Where:

`val` is the value to be written to memory.

`ptr` points to the address of the data to be written to in memory. To specify the size of the data to be written, cast the parameter to an appropriate integral type.

Table 5-12 Access widths supported by the `__strexd` intrinsic

Instruction	Size of data stored	C cast
STREXD	unsigned long long	(**unsigned long long** *)
STREXD	signed long long	(**signed long long** *)

Return value

The `__strexd` intrinsic returns:

`0` if the `STREXD` instruction succeeds

`1` if the `STREXD` instruction is locked out.

Errors

The compiler does not recognize the `__strexd` intrinsic when compiling for a target that does not support the `STREXD` instruction. The compiler generates either a warning or an error in this case.

The `__strexd` intrinsic only supports access to doubleword data. The compiler generates an error if you specify an access width that is not supported.

See also

- *__ldrex intrinsic* on page 5-68
- *__ldrexd intrinsic* on page 5-69
- *__strex intrinsic* on page 5-81
- *LDREX and STREX* on page 3-39 in the *Assembler Reference*.

5.7.38 `__strt` intrinsic

This intrinsic inserts an assembly language instruction of the form `STR{size}T` into the instruction stream generated by the compiler. It enables you to store data to memory in your C or C++ code using an `STRT` instruction.

Syntax

```
void __strt(unsigned int val, volatile void *ptr)
```

Where:

`val` Is the value to be written to memory.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

ptr Points to the address of the data to be written to in memory. To specify the size of the data to be written, cast the parameter to an appropriate integral type.

Table 5-13 Access widths supported by the __strt intrinsic

Instruction	Size of data loaded	C cast
STRBT	unsigned byte	(**char** *)
STRHT	unsigned halfword	(**short int** *)
STRT	word	(**int** *)

Errors

The compiler does not recognize the `__strt` intrinsic when compiling for a target that does not support the `STRT` instruction. The compiler generates either a warning or an error in this case.

The `__strt` intrinsic does not support access either to signed data or to doubleword data. The compiler generates an error if you specify an access width that is not supported.

Example

```
void foo(void)
{
    int loc=0xff;
    __strt(0x20, (volatile char *)loc);
}
```

Compiling this code produces:

```
||foo|| PROC
    MOV      r0,#0xff
    MOV      r1,#0x20
    STRBT    r1,[r0],#0
    BX       lr
    ENDP
```

See also

- *--thumb* on page 3-90
- *LDR and STR, unprivileged* on page 3-17 in the *Assembler Reference*.

5.7.39 `__swp` intrinsic

This intrinsic inserts a `SWP{size}` instruction into the instruction stream generated by the compiler. It enables you to swap data between memory locations from your C or C++ code.

Note

The use of `SWP` and `SWPB` is deprecated in ARMv6 and above.

Syntax

```
unsigned int __swp(unsigned int val, volatile void *ptr)
```

where:

val Is the data value to be written to memory.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

`ptr` Points to the address of the data to be written to in memory. To specify the size of the data to be written, cast the parameter to an appropriate integral type.

Table 5-14 Access widths supported by the __swp intrinsic

Instruction	Size of data loaded	C cast
SWPB	unsigned byte	(char *)
SWP	word	(int *)

Return value

The `__swp` intrinsic returns the data value that previously, is in the memory address pointed to by *ptr*, before this value is overwritten by *val*.

Example

```
int foo(void)
{
    int loc=0xff;
    return(__swp(0x20, (volatile int *)loc));
}
```

Compiling this code produces

```
||foo|| PROC
    MOV      r1, #0xff
    MOV      r0, #0x20
    SWP      r0, r0, [r1]
    BX       lr
    ENDP
```

See also

- *SWP and SWPB* on page 3-43 in the *Assembler Reference*.

5.7.40 `__usat` intrinsic

This intrinsic inserts a `USAT` instruction into the instruction stream generated by the compiler. It enables you to saturate an unsigned value from within your C or C++ code.

Syntax

```
int __usat(unsigned int val, unsigned int sat)
```

Where:

`val` Is the value to be saturated.

`sat` Is the bit position to saturate to.

usat must be in the range 0 to 31.

Return value

The `__usat` intrinsic returns *val* saturated to the unsigned range $0 \leq x \leq 2^{sat} - 1$.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Errors

The compiler does not recognize the `__usat` intrinsic when compiling for a target that does not support the `USAT` instruction. The compiler generates either a warning or an error in this case.

See also

- *__ssat intrinsic* on page 5-81
- *SSAT and USAT* on page 3-99 in the *Assembler Reference*.

5.7.41 `__wfe` intrinsic

This intrinsic inserts a `WFE` instruction into the instruction stream generated by the compiler.

In some architectures, for example the v6T2 architecture, the `WFE` instruction executes as a `NOP` instruction.

Syntax

```
void __wfe(void)
```

Errors

The compiler does not recognize the `__wfe` intrinsic when compiling for a target that does not support the `WFE` instruction. The compiler generates either a warning or an error in this case.

See also

- *__wfi intrinsic*
- *__nop* on page 5-72
- *__sev intrinsic* on page 5-79
- *__yield intrinsic* on page 5-87
- *NOP* on page 3-143 in the *Assembler Reference*.

5.7.42 `__wfi` intrinsic

This intrinsic inserts a `WFI` instruction into the instruction stream generated by the compiler.

In some architectures, for example the v6T2 architecture, the `WFI` instruction executes as a `NOP` instruction.

Syntax

```
void __wfi(void)
```

Errors

The compiler does not recognize the `__wfi` intrinsic when compiling for a target that does not support the `WFI` instruction. The compiler generates either a warning or an error in this case.

See also

- *__yield intrinsic* on page 5-87
- *__nop* on page 5-72
- *__sev intrinsic* on page 5-79
- *__wfe intrinsic*
- *NOP* on page 3-143 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.7.43 `__yield` intrinsic

This intrinsic inserts a `YIELD` instruction into the instruction stream generated by the compiler.

In some architectures, for example the v6T2 architecture, the `YIELD` instruction executes as a `NOP` instruction.

Syntax

```
void __yield(void)
```

Errors

The compiler does not recognize the `__yield` intrinsic when compiling for a target that does not support the `YIELD` instruction. The compiler generates either a warning or an error in this case.

See also

- *__nop* on page 5-72
- *__sev intrinsic* on page 5-79
- *__wfe intrinsic* on page 5-86
- *__wfi intrinsic* on page 5-86
- *NOP* on page 3-143 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.8 ARMv6 SIMD intrinsics

The ARM Architecture v6 Instruction Set Architecture adds many *Single Instruction Multiple Data* (SIMD) instructions to ARMv6 for the efficient software implementation of high-performance media applications.

The ARM compiler supports intrinsics that map to the ARMv6 SIMD instructions. These intrinsics are available when compiling your code for an ARMv6 architecture or processor. If the chosen architecture does not support the ARMv6 SIMD instructions, compilation generates a warning and subsequent linkage fails with an undefined symbol reference.

Note

Each ARMv6 SIMD intrinsic is guaranteed to be compiled into a single, inline, machine instruction for an ARM v6 architecture or processor. However, the compiler might use optimized forms of underlying instructions when it detects opportunities to do so.

The ARMv6 SIMD instructions can set the `GE[3:0]` bits in the *Application Program Status Register* (APSR). Some SIMD instructions update these flags to indicate the *greater than or equal to* status of each 8 or 16-bit slice of an SIMD operation.

The ARM compiler treats the `GE[3:0]` bits as a global variable. To access these bits from within your C or C++ program, either:

- access bits 16-19 of the APSR through a named register variable
- use the `__sel` intrinsic to control a `SEL` instruction.

5.8.1 See also

Reference

- Appendix A *ARMv6 SIMD Instruction Intrinsics* on page A-1
- *Named register variables* on page 5-94
- *ARM registers* on page 3-8 in *Using the Assembler*
- *SEL* on page 3-67 in the *Assembler Reference*
- Chapter 9 *VFP Programming* in the *Using the Assembler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.9 ETSI basic operations

The compilation tools support the original ETSI family of basic operations described in the ETSI G.729 recommendation *Coding of speech at 8 kbit/s using conjugate-structure algebraic-code-excited linear prediction (CS-ACELP)*.

To make use of the ETSI basic operations in your own code, include the standard header file `dspfns.h`. The intrinsics supplied in `dspfns.h` are listed in Table 5-15.

Table 5-15 ETSI basic operations supported by the ARM compilation tools

Intrinsics				
abs_s	L_add_c	L_mult	L_sub_c	norm_l
add	L_deposit_h	L_negate	mac_r	round
div_s	L_deposit_l	L_sat	msu_r	saturate
extract_h	L_mac	L_shl	mult	shl
extract_l	L_macNs	L_shr	mult_r	shr
L_abs	L_msu	L_shr_r	negate	shr_r
L_add	L_msuNs	L_sub	norm_s	sub

The header file `dspfns.h` also exposes certain status flags as global variables for use in your C or C++ programs. The status flags exposed by `dspfns.h` are listed in Table 5-16.

Table 5-16 ETSI status flags exposed in the ARM compilation tools

Status flag	Description
Overflow	Overflow status flag. Generally, saturating functions have a sticky effect on overflow.
Carry	Carry status flag.

5.9.1 Example

```
#include <limits.h>
#include <stdint.h>
#include <dspfns.h>       // include ETSI basic operations
int32_t C_L_add(int32_t a, int32_t b)
{
    int32_t c = a + b;
    if (((a ^ b) & INT_MIN) == 0)
    {
        if ((c ^ a) & INT_MIN)
        {
            c = (a < 0) ? INT_MIN : INT_MAX;
        }
    }
    return c;
}
__asm int32_t asm_L_add(int32_t a, int32_t b)
{
    qadd r0, r0, r1
    bx lr
}
int32_t foo(int32_t a, int32_t b)
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

```
{
    int32_t c, d, e, f;
    Overflow = 0;           // set global overflow flag
    c = C_L_add(a, b);      // C saturating add
    d = asm_L_add(a, b);    // assembly language saturating add
    e = __qadd(a, b);       // ARM intrinsic saturating add
    f = L_add(a, b);        // ETSI saturating add
    return Overflow ? -1 : c == d == e == f; // returns 1, unless overflow
}
```

5.9.2 See also

- the header file `dspfns.h` for definitions of the ETSI basic operations as a combination of C code and intrinsics
- *European Telecommunications Standards Institute (ETSI) basic operations* on page 4-12 in *Using the Compiler*
- ETSI Recommendation G.191: *Software tools for speech and audio coding standardization*
- *ITU-T Software Tool Library 2005 User's manual*, included as part of ETSI Recommendation G.191
- ETSI Recommendation G723.1 : *Dual rate speech coder for multimedia communications transmitting at 5.3 and 6.3 kbit/s*
- ETSI Recommendation G.729: *Coding of speech at 8 kbit/s using conjugate-structure algebraic-code-excited linear prediction (CS-ACELP)*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.10 C55x intrinsics

The ARM compiler supports the emulation of selected TI C55x compiler intrinsics.

To make use of the TI C55x intrinsics in your own code, include the standard header file `c55x.h`. The intrinsics supplied in `c55x.h` are listed in Table 5-17.

Table 5-17 TI C55x intrinsics supported by the compilation tools

Intrinsics			
_a_lsadd	_a_sadd	_a_smac	_a_smacr
_a_smas	_a_smasr	_abss	_count
_divs	_labss	_lmax	_lmin
_lmpy	_lmpysu	_lmpyu	_lnorm
_lsadd	_lsat	_lshl	_shrs
_lsmpy	_lsmpyi	_lsmpyr	_lsmpysu
_lsmpysui	_lsmpyu	_lsmpyui	_lsneg
_lsshl	_lssub	_max	_min
_norm	_rnd	_round	_roundn
_sadd	_shl	_shrs	_smac
_smaci	_smacr	_smacsu	_smacsui
_smas	_smasi	_smasr	_smassu
_smassui	_smpy	_sneg	_sround
_sroundn	_sshl	_ssub	-

5.10.1 Example

```
#include <limits.h>
#include <stdint.h>
#include <c55x.h>       // include TI C55x intrinsics
__asm int32_t asm_lsadd(int32_t a, int32_t b)
{
    qadd r0, r0, r1
    bx lr}
int32_t foo(int32_t a, int32_t b)
{
    int32_t c, d, e;
    c = asm_lsadd(a, b);  // assembly language saturating add
    d = __qadd(a, b);     // ARM intrinsic saturating add
    e = _lsadd(a, b);     // TI C55x saturating add
    return c == d == e;   // returns 1
}
```

5.10.2 See also

- the header file `c55x.h` for more information on the ARM implementation of the C55x intrinsics
- Publications providing information about TI compiler intrinsics are available from Texas Instruments at `http://www.ti.com`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.11 VFP status intrinsic

The compiler provides an intrinsic for reading the *Floating Point and Status Control Register* (FPSCR).

Note

ARM recommends using a named register variable as an alternative method of reading this register. This provides a more efficient method of access than using the intrinsic. See *Named register variables* on page 5-94.

5.11.1 `__vfp_status` intrinsic

This intrinsic reads or modifies the FPSCR.

Syntax

```
unsigned int __vfp_status(unsigned int mask, unsigned int flags);
```

Usage

Use this intrinsic to read or modify the flags in FPSCR.

The intrinsic returns the value of FPSCR, unmodified, if *mask* and *flags* are 0.

You can clear, set, or toggle individual flags in FPSCR using the bits in *mask* and *flags*, as shown in Table 5-18.The intrinsic returns the modified value of FPSCR if *mask* and *flags* are not both 0.

Table 5-18 Modifying the FPSCR flags

***mask* bit**	***flags* bit**	**Effect on FPSCR flag**
0	0	Does not modify the flag
0	1	Toggles the flag
1	1	Sets the flag
1	0	Clears the flag

Note

If you want to read or modify only the exception flags in FPSCR, then ARM recommends that you use the standard C99 features in `<fenv.h>`.

Errors

The compiler generates an error if you attempt to use this intrinsic when compiling for a target that does not have VFP.

See also

- *FPSCR, the floating-point status and control register* on page 9-16 in Using the Assembler
- *<fenv.h> floating-point environment access in C99 on page 5-97* in Using the Compiler.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.12 *Fused Multiply Add* (FMA) intrinsics

These intrinsics perform the following calculation, incurring only a single rounding step:

`result = a × b + c`

Performing the calculation with a single rounding step, rather than multiplying and then adding with two roundings, can result in a better degree of accuracy.

Declared in `math.h`, the FMA intrinsics are:

```
double fma(double a, double b, double c);
float fmaf(float a, float b, float c);
long double fmal(long double a, long double b, long double c);
```

Note

- These intrinsics are only available in C99 mode.
- They are only supported for the Cortex-M4 processors.
- If compiling for the Cortex-M4 processor, only `fmaf()` is available.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.13 Named register variables

The compiler enables you to access registers of an ARM architecture-based processor or coprocessor using named register variables.

5.13.1 Syntax

```
register type var-name __asm(reg);
```

Where:

type is the type of the named register variable.

Any type of the same size as the register being named can be used in the declaration of a named register variable. The type can be a structure, but bitfield layout is sensitive to endianness.

var-name is the name of the named register variable.

reg is a character string denoting the name of a register on an ARM architecture-based processor, or for coprocessor registers, a string syntax that identifies the coprocessor and corresponds with how you intend to use the variable.

Registers available for use with named register variables on ARM architecture-based processors are shown in Table 5-19.

Table 5-19 Named registers available on ARM architecture-based processors

Register	Character string for `__asm`	Processors
`APSR`	`"apsr"`	All processors
`CPSR`	`"cpsr"`	All processors
`BASEPRI`	`"basepri"`	Cortex-M3, Cortex-M4
`BASEPRI_MAX`	`"basepri_max"`	Cortex-M3, Cortex-M4
`CONTROL`	`"control"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4
`DSP`	`"dsp"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4
`EAPSR`	`"eapsr"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4
`EPSR`	`"epsr"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4
`FAULTMASK`	`"faultmask"`	Cortex-M3, Cortex-M4
`IAPSR`	`"iapsr"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4
`IEPSR`	`"iepsr"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4
`IPSR`	`"ipsr"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4
`MSP`	`"msp"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table 5-19 Named registers available on ARM architecture-based processors (continued)

Register	Character string for `__asm`	Processors
`PRIMASK`	`"primask"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4
`PSP`	`"psp"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4
`PSR`	`"psr"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4
`r0` to `r12`	`"r0"` to `"r12"`	All processors
`r13` or `sp`	`"r13"` or `"sp"`	All processors
`r14` or `lr`	`"r14"` or `"lr"`	All processors
`r15` or `pc`	`"r15"` or `"pc"`	All processors
`SPSR`	`"spsr"`	All processors, apart from Cortex-M series processors.
`XPSR`	`"xpsr"`	Cortex-M0, Cortex-M1, Cortex-M3, Cortex-M4

On targets with floating-point hardware, the registers of Table 5-20 are also available for use with named register variables.

Table 5-20 Named registers available on targets with floating-point hardware

Register	Character string for __asm
`FPSID`	`"fpsid"`
`FPSCR`	`"fpscr"`
`FPEXC`	`"fpexc"`
`FPINST`	`"fpinst"`
`FPINST2`	`"fpinst2"`
`FPSR`	`"fpsr"`
`MVFR0`	`"mvfr0"`
`MVFR1`	`"mvfr1"`

Note

Some registers are not available on some architectures.

5.13.2 Usage

You can declare named register variables as global variables. You can declare some, but not all, named register variables as local variables. In general, do not declare VFP registers and core registers as local variables. Do not declare caller-save registers, such as `R0`, as local variables.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.13.3 Examples

In Example 5-8, `apsr` is declared as a named register variable for the `"apsr"` register:

Example 5-8 Named register variable for APSR

```
register unsigned int apsr __asm("apsr");
apsr = ~(~apsr | 0x40);
```

This generates the following instruction sequence:

```
MRS r0,APSR ; formerly CPSR
BIC r0,r0,#0x40
MSR CPSR_c, r0
```

In Example 5-9, `PMCR` is declared as a register variable associated with coprocessor `cp15`, with `CRn = c9`, `CRm = c12`, `opcode1 = 0`, and `opcode2 = 0`, in an `MCR` or an `MRC` instruction:

Example 5-9 Named register variable for coprocessor register

```
register unsigned int PMCR __asm("cp15:0:c9:c12:0");

__inline void __reset_cycle_counter(void)
{
    PMCR = 4;
}
```

The disassembled output is as follows:

```
__reset_cycle_counter PROC
    MOV     r0,#4
    MCR     p15,#0x0,r0,c9,c12,#0
    BX      lr
    ENDP
```

In Example 5-10, `cp15_control` is declared as a register variable that is used to access a coprocessor register. This example enables the MMU using `CP15`:

Example 5-10 Named register variable for coprocessor register to enable MMU

```
register unsigned int cp15_control __asm("cp15:0:c1:c0:0");
cp15_control |= 0x1;
```

The following instruction sequence is generated:

```
MRC  p15,#0x0,r0,c1,c0,#0
ORR  r0,r0,#1
MCR  p15,#0x0,r0,c1,c0,#0
```

Example 5-11 on page 5-97 for Cortex-M3 declares `_msp`, `_control` and `_psp` as named register variables to set up stack pointers:

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Example 5-11 Named register variables to set up stack pointers on Cortex-M3

```
register unsigned int _control __asm("control");
register unsigned int _msp     __asm("msp");
register unsigned int _psp     __asm("psp");void init(void)
{
  _msp = 0x30000000;        // set up Main Stack Pointer
  _control = _control | 3;  // switch to User Mode with Process Stack
  _psp = 0x40000000;        // setup Process Stack Pointer
}
```

This generates the following instruction sequence:

```
init
  MOV r0,#0x30000000
  MSR MSP,r0
  MRS r0,CONTROL
  ORR r0,r0,#3
  MSR CONTROL,r0
  MOV r0,#0x40000000
  MSR PSP,r0
  BX lr
```

5.13.4 See also

- *Compiler support for accessing registers using named register variables* on page 4-16 in *Using the Compiler*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

5.14 Compiler predefines

This section documents the predefined macros of the ARM compiler.

5.14.1 Predefined macros

Table 5-21 lists the macro names predefined by the ARM compiler for C and C++. Where the value field is empty, the symbol is only defined.

Table 5-21 Predefined macros

Name	Value	When defined
`__arm__`	–	Always defined for the ARM compiler, even when you specify the `--thumb` option. See also `__ARMCC_VERSION`.
`__ARMCC_VERSION`	*ver*	Always defined. It is a decimal number, and is guaranteed to increase between releases. The format is *PVbbbb* where: • *P* is the major version • *V* is the minor version • *bbbb* is the build number. **Note** Use this macro to distinguish between ARM Compiler 4.1 or later, and other tools that define `__arm__`.
`__APCS_INTERWORK`	–	When you specify the `--apcs /interwork` option or set the CPU architecture to ARMv5T or later.
`__APCS_ROPI`	–	When you specify the `--apcs /ropi` option.
`__APCS_RWPI`	–	When you specify the `--apcs /rwpi` option.
`__APCS_FPIC`	–	When you specify the `--apcs /fpic` option.
`__ARRAY_OPERATORS`	–	In C++ compiler mode, to specify that array new and delete are enabled.
`__BASE_FILE__`	*name*	Always defined. Similar to `__FILE__`, but indicates the primary source file rather than the current one (that is, when the current file is an included file).
`__BIG_ENDIAN`	–	If compiling for a big-endian target.
`_BOOL`	–	In C++ compiler mode, to specify that **`bool`** is a keyword.
`__cplusplus`	–	In C++ compiler mode.
`__CC_ARM`	1	Always set to 1 for the ARM compiler, even when you specify the `--thumb` option.
`__DATE__`	*date*	Always defined.
`__EDG__`	–	Always defined.
`__EDG_IMPLICIT_USING_STD`	–	In C++ mode when you specify the `--using_std` option.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table 5-21 Predefined macros (continued)

Name	Value	When defined
`__EDG_VERSION__`	–	Always set to an integer value that represents the version number of the *Edison Design Group* (EDG) front-end. For example, version 3.8 is represented as `308`. *The version number of the EDG front-end does not necessarily match the version number of the ARM compiler toolchain.*
`__EXCEPTIONS`	1	In C++ mode when you specify the `--exceptions` option.
`__FEATURE_SIGNED_CHAR`	–	When you specify the `--signed_chars` option (used by `CHAR_MIN` and `CHAR_MAX`).
`__FILE__`	*name*	Always defined as a string literal.
`__FP_FAST`	–	When you specify the `--fpmode=fast` option.
`__FP_FENV_EXCEPTIONS`	–	When you specify the `--fpmode=ieee_full` or `--fpmode=ieee_fixed` options.
`__FP_FENV_ROUNDING`	–	When you specify the `--fpmode=ieee_full` option.
`__FP_IEEE`	–	When you specify the `--fpmode=ieee_full`, `--fpmode=ieee_fixed`, or `--fpmode=ieee_no_fenv` options.
`__FP_INEXACT_EXCEPTION`	–	When you specify the `--fpmode=ieee_full` option.
`__IMPLICIT_INCLUDE`	–	When you specify the `--implicit_include` option.
`__LINE__`	*num*	Always set. It is the source line number of the line of code containing this macro.
`__MODULE__`	*mod*	Contains the filename part of the value of `__FILE__`.
`__MULTIFILE`	–	When you explicitly or implicitly use the `--multifile` option.[a]
`__OPTIMISE_LEVEL`	*num*	Always set to 2 by default, unless you change the optimization level using the `-O`*num* option.[a]
`__OPTIMISE_SPACE`	–	When you specify the `-Ospace` option.
`__OPTIMISE_TIME`	–	When you specify the `-Otime` option.
`__PLACEMENT_DELETE`	–	In C++ mode to specify that placement delete (that is, an operator **`delete`** corresponding to a placement operator **`new`**, to be called if the constructor throws an exception) is enabled. This is only relevant when using exceptions.
`__RTTI`	–	In C++ mode when RTTI is enabled.
`__sizeof_int`	4	For `sizeof(int)`, but available in preprocessor expressions.
`__sizeof_long`	4	For `sizeof(long)`, but available in preprocessor expressions.
`__sizeof_ptr`	4	For `sizeof(void *)`, but available in preprocessor expressions.
`__SOFTFP__`	–	If compiling to use the software floating-point calling standard and library. Set when you specify the `--fpu=softvfp` option for ARM or Thumb, or when you specify `--fpu=softvfp+vfpv2` for Thumb.
`__STDC__`	–	In all compiler modes.
`__STDC_VERSION__`	–	Standard version information.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table 5-21 Predefined macros (continued)

Name	Value	When defined
`__STRICT_ANSI__`	–	When you specify the `--strict` option.
`__SUPPORT_SNAN__`	–	Support for signalling NaNs when you specify `--fpmode=ieee_fixed` or `--fpmode=ieee_full`.
`__TARGET_ARCH_ARM`	*num*	The number of the ARM base architecture of the target CPU irrespective of whether the compiler is compiling for ARM or Thumb. For possible values of `__TARGET_ARCH_ARM` in relation to the ARM architecture versions, see Table 5-22 on page 5-102.
`__TARGET_ARCH_THUMB`	*num*	The number of the Thumb base architecture of the target CPU irrespective of whether the compiler is compiling for ARM or Thumb. The value is defined as zero if the target does not support Thumb. For possible values of `__TARGET_ARCH_THUMB` in relation to the ARM architecture versions, see Table 5-22 on page 5-102.
`__TARGET_ARCH_`*`XX`*	–	*`XX`* represents the target architecture and its value depends on the target architecture. For example, if you specify the compiler options `--cpu=4T` or `--cpu=ARM7TDMI` then `__TARGET_ARCH_4T` is defined.
`__TARGET_CPU_`*`XX`*	–	*`XX`* represents the target CPU. The value of *`XX`* is derived from the `--cpu` compiler option, or the default if none is specified. For example, if you specify the compiler option `--cpu=ARM7TM` then `__TARGET_CPU_ARM7TM` is defined and no other symbol starting with `__TARGET_CPU_` is defined. If you specify the target architecture, then `__TARGET_CPU_generic` is defined. If the CPU name specified with `--cpu` is in lowercase, it is converted to uppercase. For example, `--cpu=Cortex-R4` results in `__TARGET_CPU_CORTEX_R4` being defined (rather than `__TARGET_CPU_Cortex_R4`). If the processor name contains hyphen (-) characters, these are mapped to an underscore (_). For example, `--cpu=ARM926EJ-S` is mapped to `__TARGET_CPU_ARM926EJ_S`.
`__TARGET_FEATURE_DOUBLEWORD`	–	ARMv5T and above.
`__TARGET_FEATURE_DSPMUL`	–	If the DSP-enhanced multiplier is available, for example ARMv5TE.
`__TARGET_FEATURE_MULTIPLY`	–	If the target architecture supports the long multiply instructions `MULL` and `MULAL`.
`__TARGET_FEATURE_DIVIDE`	–	If the target architecture supports the hardware divide instruction (that is, ARMv7-M or ARMv7-R).
`__TARGET_FEATURE_THUMB`	–	If the target architecture supports Thumb, ARMv4T or later.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table 5-21 Predefined macros (continued)

Name	Value	When defined
`__TARGET_FPU_xx`	–	One of the following is set to indicate the FPU usage: • `__TARGET_FPU_NONE` • `__TARGET_FPU_VFP` • `__TARGET_FPU_SOFTVFP` In addition, if compiling with one of the following `--fpu` options, the corresponding target name is set: • `--fpu=softvfp+vfpv2`, `__TARGET_FPU_SOFTVFP_VFPV2` • `--fpu=softvfp+vfpv3`, `__TARGET_FPU_SOFTVFP_VFPV3` • `--fpu=softvfp+vfpv3_fp16`, `__TARGET_FPU_SOFTVFP_VFPV3_FP16` • `--fpu=softvfp+vfpv3_d16`, `__TARGET_FPU_SOFTVFP_VFPV3_D16` • `--fpu=softvfp+vfpv3_d16_fp16`, `__TARGET_FPU_SOFTVFP_VFPV3_D16_FP16` • `--fpu=vfpv2`, `__TARGET_FPU_VFPV2` • `--fpu=vfpv3`, `__TARGET_FPU_VFPV3` • `--fpu=vfpv3_fp16`, `__TARGET_FPU_VFPV3_FP16` • `--fpu=vfpv3_d16`, `__TARGET_FPU_VFPV3_D16` • `--fpu=vfpv3_d16_fp16`, `__TARGET_FPU_VFPV3_D16_FP16` See *--fpu=name* on page 3-44 for more information.
`__TARGET_PROFILE_A`		When you specify the `--cpu=7-A` option.
`__TARGET_PROFILE_R`		When you specify the `--cpu=7-R` option.
`__TARGET_PROFILE_M`		When you specify any of the following options: • `--cpu=6-M` • `--cpu=6S-M` • `--cpu=7-M`
`__thumb__`	–	When the compiler is in Thumb mode. That is, you have either specified the `--thumb` option on the command-line or `#pragma thumb` in your source code. **Note** • The compiler might generate some ARM code even if it is compiling for Thumb. • `__thumb` and `__thumb__` become defined or undefined when using `#pragma thumb` or `#pragma arm`, but do not change in cases where Thumb functions are generated as ARM code for other reasons (for example, a function specified as `__irq`).
`__TIME__`	*time*	Always defined.
`_WCHAR_T`	–	In C++ mode, to specify that **`wchar_t`** is a keyword.

a. ARM recommends that if you have source code reliant on the `__OPTIMISE_LEVEL` macro to determine whether or not `--multifile` is in effect, you change to using `__MULTIFILE`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table 5-22 shows the possible values for __TARGET_ARCH_THUMB (see Table 5-21 on page 5-98), and how these values relate to versions of the ARM architecture.

Table 5-22 Thumb architecture versions in relation to ARM architecture versions

ARM architecture	__TARGET_ARCH_ARM	__TARGET_ARCH_THUMB
v4	4	0
v4T	4	1
v5T, v5TE, v5TEJ	5	2
v6, v6K, v6Z	6	3
v6T2	6	4
v6-M, v6S-M	0	3
v7-A, v7-R	7	4
v7-M, v7E-M	0	4

5.14.2 Function names

Table 5-23 lists builtin variables supported by the compiler for C and C++.

Table 5-23 Builtin variables

Name	Value
__FUNCTION__	Holds the name of the function as it appears in the source. __FUNCTION__ is a constant string literal. You cannot use the preprocessor to join the contents to other text to form new tokens.
__PRETTY_FUNCTION__	Holds the name of the function as it appears pretty printed in a language-specific fashion. __PRETTY_FUNCTION__ is a constant string literal. You cannot use the preprocessor to join the contents to other text to form new tokens.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Chapter 6
C and C++ Implementation Details

This chapter describes the language implementation details for the compiler. It includes:

- *C and C++ implementation details* on page 6-2
- *C++ implementation details* on page 6-11.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

6.1 C and C++ implementation details

This section describes language implementation details common to both C and C++.

6.1.1 Character sets and identifiers

The following points apply to the character sets and identifiers expected by the compiler:

- Uppercase and lowercase characters are distinct in all internal and external identifiers. An identifier can also contain a dollar (`$`) character unless the `--strict` compiler option is specified. To permit dollar signs in identifiers with the `--strict` option, also use the `--dollar` command-line option.
- Calling `setlocale(LC_CTYPE, "ISO8859-1")` makes the `isupper()` and `islower()` functions behave as expected over the full 8-bit Latin-1 alphabet, rather than over the 7-bit ASCII subset. The locale must be selected at link time.
- Source files are compiled according to the currently selected locale. You might have to select a different locale, with the `--locale` command-line option, if the source file contains non-ASCII characters. See *Compiler command-line options listed by group* on page 3-4 in *Using the Compiler* for more information.
- The compiler supports multibyte character sets, such as Unicode.
- Other properties of the source character set are host-specific.

The properties of the execution character set are target-specific. The ARM C and C++ libraries support the ISO 8859-1 (Latin-1 Alphabet) character set with the following consequences:

- The execution character set is identical to the source character set.
- There are eight bits in a character in the execution character set.
- There are four characters (bytes) in an **`int`**. If the memory system is:

  Little-endian The bytes are ordered from least significant at the lowest address to most significant at the highest address.

  Big-endian The bytes are ordered from least significant at the highest address to most significant at the lowest address.

- In C all character constants have type **`int`**. In C++ a character constant containing one character has the type **`char`** and a character constant containing more than one character has the type **`int`**. Up to four characters of the constant are represented in the integer value. The last character in the constant occupies the lowest-order byte of the integer value. Up to three preceding characters are placed at higher-order bytes. Unused bytes are filled with the `NULL` (`\0`) character.
- Table 6-1 lists all integer character constants, that contain a single character or character escape sequence, are represented in both the source and execution character sets.

Table 6-1 Character escape codes

Escape sequence	Char value	Description
`\a`	`7`	Attention (bell)
`\b`	`8`	Backspace
`\t`	`9`	Horizontal tab
`\n`	`10`	New line (line feed)

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table 6-1 Character escape codes (continued)

Escape sequence	Char value	Description
`\v`	`11`	Vertical tab
`\f`	`12`	Form feed
`\r`	`13`	Carriage return
`\xnn`	`0xnn`	ASCII code in hexadecimal
`\nnn`	`0nnn`	ASCII code in octal

- Characters of the source character set in string literals and character constants map identically into the execution character set.
- Data items of type **`char`** are unsigned by default. They can be explicitly declared as **`signed char`** or **`unsigned char`**:
  - — the `--signed_chars` option can be used to make the **`char`** signed
  - — the `--unsigned_chars` option can be used to make the **`char`** unsigned.

  Note

  Care must be taken when mixing translation units that have been compiled with and without the `--signed_chars` and `--unsigned_chars` options, and that share interfaces or data structures.

  The ARM ABI defines **`char`** as an unsigned byte, and this is the interpretation used by the C++ libraries supplied with the ARM compilation tools.

- No locale is used to convert multibyte characters into the corresponding wide characters for a wide character constant. This is not relevant to the generic implementation.

6.1.2 Basic data types

This section describes how the basic data types are implemented in ARM C and C++.

Size and alignment of basic data types

Table 6-2 gives the size and natural alignment of the basic data types.

Table 6-2 Size and alignment of data types

Type	Size in bits	Natural alignment in bytes
`char`	8	1 (byte-aligned)
`short`	16	2 (halfword-aligned)
`int`	32	4 (word-aligned)
`long`	32	4 (word-aligned)
`long long`	64	8 (doubleword-aligned)
`float`	32	4 (word-aligned)
`double`	64	8 (doubleword-aligned)
`long double`	64	8 (doubleword-aligned)

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table 6-2 Size and alignment of data types (continued)

Type	Size in bits	Natural alignment in bytes
All pointers	32	4 (word-aligned)
`bool` (C++ only)	8	1 (byte-aligned)
`_Bool` (C only[a])	8	1 (byte-aligned)
`wchar_t` (C++ only)	16	2 (halfword-aligned)

a. `stdbool.h` can be used to define the `bool` macro in C.

Type alignment varies according to the context:

- Local variables are usually kept in registers, but when local variables spill onto the stack, they are always word-aligned. For example, a spilled local **`char`** variable has an alignment of 4.
- The natural alignment of a packed type is 1.

See *Structures, unions, enumerations, and bitfields* on page 6-6 for more information.

Integer

Integers are represented in two's complement form. The low word of a **`long long`** is at the low address in little-endian mode, and at the high address in big-endian mode.

Float

Floating-point quantities are stored in IEEE format:
- **`float`** values are represented by IEEE single-precision values
- **`double`** and **`long double`** values are represented by IEEE double-precision values.

For **`double`** and **`long double`** quantities the word containing the sign, the exponent, and the most significant part of the mantissa is stored with the lower machine address in big-endian mode and at the higher address in little-endian mode. See *Operations on floating-point types* on page 6-5 for more information.

Arrays and pointers

The following statements apply to all pointers to objects in C and C++, except pointers to members:
- Adjacent bytes have addresses that differ by one.
- The macro `NULL` expands to the value 0.
- Casting between integers and pointers results in no change of representation.
- The compiler warns of casts between pointers to functions and pointers to data.
- The type `size_t` is defined as `unsigned int`.
- The type `ptrdiff_t` is defined as `signed int`.

6.1.3 Operations on basic data types

The ARM compiler performs the usual arithmetic conversions set out in relevant sections of the ISO C99 and ISO C++ standards. The following subsections describe additional points that relate to arithmetic operations.

See also *Expression evaluation* on page D-6.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Operations on integral types

The following statements apply to operations on the integral types:

- All signed integer arithmetic uses a two's complement representation.
- Bitwise operations on signed integral types follow the rules that arise naturally from two's complement representation. No sign extension takes place.
- Right shifts on signed quantities are arithmetic.
- For values of type `int`,
  - Shifts outside the range 0 to 127 are undefined.
  - Left shifts of more than 31 give a result of zero.
  - Right shifts of more than 31 give a result of zero from a shift of an unsigned value or positive signed value. They yield –1 from a shift of a negative signed value.
- For values of type `long long`, shifts outside the range 0 to 63 are undefined.
- The remainder on integer division has the same sign as the numerator, as mandated by the ISO C99 standard.
- If a value of integral type is truncated to a shorter signed integral type, the result is obtained by discarding an appropriate number of most significant bits. If the original number is too large, positive or negative, for the new type, there is no guarantee that the sign of the result is going to be the same as the original.
- A conversion between integral types does not raise an exception.
- Integer overflow does not raise an exception.
- Integer division by zero returns zero by default.

Operations on floating-point types

The following statements apply to operations on floating-point types:

- Normal IEEE 754 rules apply.
- Rounding is to the nearest representable value by default.
- Floating-point exceptions are disabled by default.

Also, see *--fpmode=model* on page 3-42.

Note

The IEEE 754 standard for floating-point processing states that the default action to an exception is to proceed without a trap. You can modify floating-point error handling by tailoring the functions and definitions in `fenv.h`. See *Modification of C library functions for error signaling, error handling, and program exit* on page 2-80 in *Using ARM® C and C++ Libraries and Floating-Point Support* for more information.

Pointer subtraction

The following statements apply to all pointers in C. They also apply to pointers in C++, other than pointers to members:

- When one pointer is subtracted from another, the difference is the result of the expression:

```
((int)a - (int)b) / (int)sizeof(type pointed to)
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- If the pointers point to objects whose alignment is the same as their size, this alignment ensures that division is exact.
- If the pointers point to objects whose alignment is less than their size, such as packed types and most **struct**s, both pointers must point to elements of the same array.

6.1.4 Structures, unions, enumerations, and bitfields

This section describes the implementation of the structured data types union, enum, and struct. It also discusses structure padding and bitfield implementation.

See *Anonymous classes, structures and unions* on page 4-16 for more information.

Unions

When a member of a **union** is accessed using a member of a different type, the resulting value can be predicted from the representation of the original type. No error is given.

Enumerations

An object of type **enum** is implemented in the smallest integral type that contains the range of the **enum**.

In C mode, and in C++ mode without `--enum_is_int`, if an **enum** contains only positive enumerator values, the storage type of the **enum** is the first *unsigned* type from the following list, according to the range of the enumerators in the **enum**. In other modes, and in cases where an **enum** contains any negative enumerator values, the storage type of the **enum** is the first of the following, according to the range of the enumerators in the **enum**:

- **unsigned char** if not using `--enum_is_int`
- **signed char** if not using `--enum_is_int`
- **unsigned short** if not using `--enum_is_int`
- **signed short** if not using `--enum_is_int`
- **signed int**
- **unsigned int** except C with `--strict`
- **signed long long** except C with `--strict`
- **unsigned long long** except C with `--strict`.

Note

- In ARM Compiler 4.1 and later, the storage type of the **enum** being the first unsigned type from the list applies irrespective of mode.

Implementing **enum** in this way can reduce data size. The command-line option `--enum_is_int` forces the underlying type of **enum** to at least as wide as **int**.

See the description of C language mappings in the *Procedure Call Standard for the ARM Architecture* specification for more information.

Note

Care must be taken when mixing translation units that have been compiled with and without the `--enum_is_int` option, and that share interfaces or data structures.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Handling values that are out of range

In strict C, enumerator values must be representable as **int**s, for example, they must be in the range -2147483648 to +2147483647, inclusive. In some earlier releases of RVCT out-of-range values were cast to **int** without a warning (unless you specified the `--strict` option).

In RVCT v2.2 and later, a Warning is issued for out-of-range enumerator values:

```
#66: enumeration value is out of "int" range
```

Such values are treated the same way as in C++, that is, they are treated as **unsigned int**, **long long**, or **unsigned long long**.

To ensure that out-of-range Warnings are reported, use the following command to change them into Errors:

```
armcc --diag_error=66 ...
```

Structures

The following points apply to:

- all C structures
- all C++ structures and classes not using virtual functions or base classes.

Structure alignment

The alignment of a nonpacked structure is the maximum alignment required by any of its fields.

Field alignment

Structures are arranged with the first-named component at the lowest address. Fields are aligned as follows:

- A field with a **char** type is aligned to the next available byte.
- A field with a **short** type is aligned to the next even-addressed byte.
- In RVCT v2.0 and above, **double** and **long long** data types are eight-byte aligned. This enables efficient use of the `LDRD` and `STRD` instructions in ARMv5TE and above.
- Bitfield alignment depends on how the bitfield is declared. See *Bitfields in packed structures* on page 6-10 for more information.
- All other types are aligned on word boundaries.

Structures can contain padding to ensure that fields are correctly aligned and that the structure itself is correctly aligned. Figure 6-1 on page 6-8 shows an example of a conventional, nonpacked structure. Bytes 1, 2, and 3 are padded to ensure correct field alignment. Bytes 11 and 12 are padded to ensure correct structure alignment. The `sizeof()` function returns the size of the structure including padding.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

```
struct {char c; int x; short s} ex1;
```

0	1	2	3
c	padding		
4	5	7	8
x			
9	10	11	12
s		padding	

Figure 6-1 Conventional nonpacked structure example

The compiler pads structures in one of the following ways, according to how the structure is defined:

- Structures that are defined as **`static`** or **`extern`** are padded with zeros.
- Structures on the stack or heap, such as those defined with `malloc()` or **`auto`**, are padded with whatever is previously stored in those memory locations. You cannot use `memcmp()` to compare padded structures defined in this way (see Figure 6-1).

Use the `--remarks` option to view the messages that are generated when the compiler inserts padding in a **`struct`**.

Structures with empty initializers are permitted in C++:

```
struct
{
    int x;
} X = { };
```

However, if you are compiling C, or compiling C++ with the `--cpp` and`--c90` options, an error is generated.

Packed structures

A packed structure is one where the alignment of the structure, and of the fields within it, is always 1.

You can pack specific structures with the `__packed` qualifier. Alternatively, you can use `#pragma pack(`*`n`*`)` to make sure that any structures with unaligned data are packed. There is no command-line option to change the default packing of structures.

Bitfields

In nonpacked structures, the ARM compiler allocates bitfields in *containers*. A container is a correctly aligned object of a declared type.

Bitfields are allocated so that the first field specified occupies the lowest-addressed bits of the word, depending on configuration:

Little-endian Lowest addressed means least significant.

Big-endian Lowest addressed means most significant.

A bitfield container can be any of the integral types.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Note

In strict 1990 ISO Standard C, the only types permitted for a bit field are **`int`**, **`signed int`**, and **`unsigned int`**. For non-**`int`** bitfields, the compiler displays an error.

A plain bitfield, declared without either **`signed`** or **`unsigned`** qualifiers, is treated as **`unsigned`**. For example, `int x:10` allocates an unsigned integer of 10 bits.

A bitfield is allocated to the first container of the correct type that has a sufficient number of unallocated bits, for example:

```
struct X
{
    int x:10;
    int y:20;
};
```

The first declaration creates an integer container and allocates 10 bits to `x`. At the second declaration, the compiler finds the existing integer container with a sufficient number of unallocated bits, and allocates `y` in the same container as `x`.

A bitfield is wholly contained within its container. A bitfield that does not fit in a container is placed in the next container of the same type. For example, the declaration of `z` overflows the container if an additional bitfield is declared for the structure:

```
struct X
{
    int x:10;
    int y:20;
    int z:5;
};
```

The compiler pads the remaining two bits for the first container and assigns a new integer container for `z`.

Bitfield containers can *overlap* each other, for example:

```
struct X
{
    int x:10;
    char y:2;
};
```

The first declaration creates an integer container and allocates 10 bits to `x`. These 10 bits occupy the first byte and two bits of the second byte of the integer container. At the second declaration, the compiler checks for a container of type **`char`**. There is no suitable container, so the compiler allocates a new correctly aligned **`char`** container.

Because the natural alignment of **`char`** is 1, the compiler searches for the first byte that contains a sufficient number of unallocated bits to completely contain the bitfield. In the example structure, the second byte of the **`int`** container has two bits allocated to `x`, and six bits unallocated. The compiler allocates a **`char`** container starting at the second byte of the previous **`int`** container, skips the first two bits that are allocated to `x`, and allocates two bits to `y`.

If `y` is declared `char y:8`, the compiler pads the second byte and allocates a new **`char`** container to the third byte, because the bitfield cannot overflow its container. Figure 6-2 on page 6-10 shows the bitfield allocation for the following example structure:

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

```
struct X
{
    int x:10;
    char y:8;
};
```

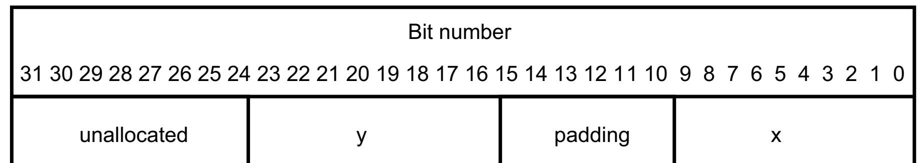

Figure 6-2 Bitfield allocation 1

Note

The same basic rules apply to bitfield declarations with different container types. For example, adding an **`int`** bitfield to the example structure gives:

```
struct X
{
    int x:10;
    char y:8;
    int z:5;
}
```

The compiler allocates an **`int`** container starting at the same location as the `int x:10` container and allocates a byte-aligned **`char`** and 5-bit bitfield, see Figure 6-3.

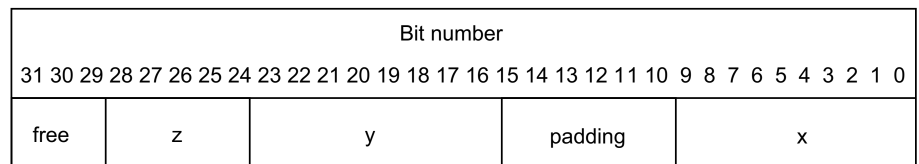

Figure 6-3 Bitfield allocation 2

You can explicitly pad a bitfield container by declaring an unnamed bitfield of size zero. A bitfield of zero size fills the container up to the end if the container is not empty. A subsequent bitfield declaration starts a new empty container.

Bitfields in packed structures

Bitfield containers in packed structures have an alignment of 1. Therefore, the maximum bit padding for a bitfield in a packed structure is 7 bits. For an unpacked structure, the maximum padding is `8*sizeof(container-type)-1 bits`.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

6.2 C++ implementation details

This section describes language implementation details specific to C++.

6.2.1 Using the ::operator new function

In accordance with the ISO C++ Standard, the `::operator new(std::size_t)` throws an exception when memory allocation fails rather than raising a signal. If the exception is not caught, `std::terminate()` is called.

The compiler option `--force_new_nothrow` turns all new calls in a compilation into calls to `::operator new(std::size_t, std::nothrow_t&)` or `:operator new[](std::size_t, std::nothrow_t&)`. However, this does not affect `operator new` calls in libraries, nor calls to any class-specific `operator new`. See *--force_new_nothrow, --no_force_new_nothrow* on page 3-40 for more information.

Legacy support

In RVCT v2.0, when the `::operator new` function ran out of memory, it raised the signal **SIGOUTOFHEAP**, instead of throwing a C++ exception. See *ISO-compliant implementation of signals supported by the signal() function in the C library and additional type arguments* on page 2-110 in *Using ARM® C and C++ Libraries and Floating-Point Support*.

In the current release, it is possible to install a `new_handler` to raise a signal and so restore the RVCT v2.0 behavior.

Note

Do not rely on the implementation details of this behavior, because it might change in future releases.

6.2.2 Tentative arrays

The ADS v1.2 and RVCT v1.2 C++ compilers enabled you to use tentative, that is, incomplete array declarations, for example, `int a[]`. You cannot use tentative arrays when compiling C++ with the RVCT v2.x compilers or later, or with ARM Compiler 4.1 or later.

6.2.3 Old-style C parameters in C++ functions

The ADS v1.2 and RVCT v1.2 C++ compilers enabled you to use old-style C parameters in C++ functions. That is,

```
void f(x) int x; { }
```

In the RVCT v2.x compilers or above, you must use the `--anachronisms` compiler option if your code contains any old-style parameters in functions. The compiler warns you if it finds any instances.

6.2.4 Anachronisms

The following anachronisms are accepted when you enable anachronisms using the `--anachronisms` option:

- **`overload`** is permitted in function declarations. It is accepted and ignored.
- Definitions are not required for static data members that can be initialized using default initialization. The anachronism does not apply to static data members of template classes, because these must always be defined.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- The number of elements in an array can be specified in an array delete operation. The value is ignored.
- A single `operator++()` and `operator--()` function can be used to overload both prefix and postfix operations.
- The base class name can be omitted in a base class initializer if there is only one immediate base class.
- Assignment to the `this` pointer in constructors and destructors is permitted.
- A bound function pointer, that is, a pointer to a member function for a given object, can be cast to a pointer to a function.
- A nested class name can be used as a non-nested class name provided no other class of that name has been declared. The anachronism is not applied to template classes.
- A reference to a non-`const` type can be initialized from a value of a different type. A temporary is created, it is initialized from the converted initial value, and the reference is set to the temporary.
- A reference to a non `const` class type can be initialized from an rvalue of the class type or a class derived from that class type. No, additional, temporary is used.
- A function with old-style parameter declarations is permitted and can participate in function overloading as if it were prototyped. Default argument promotion is not applied to parameter types of such functions when the check for compatibility is done, so that the following declares the overloading of two functions named `f`:

```
int f(int);
int f(x) char x; { return x; }
```

  Note

  In C, this code is legal but has a different meaning. A tentative declaration of `f` is followed by its definition.

6.2.5 Template instantiation

The compiler does all template instantiations automatically, and makes sure there is only one definition of each template entity left after linking. The compiler does this by emitting template entities in named common sections. Therefore, all duplicate common sections, that is, common sections with the same name, are eliminated by the linker.

Note

You can limit the number of concurrent instantiations of a given template with the `--pending_instantiations` compiler option.

See also *--pending_instantiations=n* on page 3-76 for more information.

Implicit inclusion

When implicit inclusion is enabled, the compiler assumes that if it requires a definition to instantiate a template entity declared in a `.h` file it can implicitly include the corresponding `.cc` file to get the source code for the definition. For example, if a template entity `ABC::f` is declared in file `xyz.h`, and an instantiation of `ABC::f` is required in a compilation but no definition of `ABC::f` appears in the source code processed by the compilation, then the compiler checks to see if a file `xyz.cc` exists. If this file exists, the compiler processes the file as if it were included at the end of the main source file.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

To find the template definition file for a given template entity the compiler has to know the full path name of the file where the template is declared and whether the file is included using the system include syntax, for example, `#include <file.h>`. This information is not available for preprocessed source containing `#line` directives. Consequently, the compiler does not attempt implicit inclusion for source code containing `#line` directives.

The compiler looks for the definition-file suffixes `.cc` and `.CC`.

You can turn implicit inclusion mode on or off with the command-line options `--implicit_include` and `--no_implicit_include`.

Implicit inclusions are only performed during the normal compilation of a file, that is, when not using the `-E` command-line option.

See *Command-line options* on page 3-6 for more information.

6.2.6 Namespaces

When doing name lookup in a template instantiation, some names must be found in the context of the template definition. Other names can be found in the context of the template instantiation. The compiler implements two different instantiation lookup algorithms:

- the algorithm mandated by the standard, and referred to as dependent name lookup.
- the algorithm that exists before dependent name lookup is implemented.

Dependent name lookup is done in strict mode, unless explicitly disabled by another command-line option, or when dependent name processing is enabled by either a configuration flag or a command-line option.

Dependent name lookup processing

When doing dependent name lookup, the compiler implements the instantiation name lookup rules specified in the standard. This processing requires that nonclass prototype instantiations be done. This in turn requires that the code be written using the typename and template keywords as required by the standard.

Lookup using the referencing context

When not using dependent name lookup, the compiler uses a name lookup algorithm that approximates the two-phase lookup rule of the standard, but in a way that is more compatible with existing code and existing compilers.

When a name is looked up as part of a template instantiation, but is not found in the local context of the instantiation, it is looked up in a synthesized instantiation context. This synthesized instantiation context includes both names from the context of the template definition and names from the context of the instantiation. For example:

```
namespace N
{
    int g(int);
    int x = 0;
    template <class T> struct A
    {
        T f(T t) { return g(t); }
        T f() { return x; }
    };
}
namespace M {
    int x = 99;
```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

```
    double g(double);
    N::A<int> ai;
    int i = ai.f(0);        // N::A<int>::f(int) calls N::g(int)
    int i2 = ai.f();        // N::A<int>::f() returns 0 (= N::x)
    N::A<double> ad;
    double d = ad.f(0);     // N::A<double>::f(double) calls M::g(double)
    double d2 = ad.f();     // N::A<double>::f() also returns 0 (= N::x)
}
```

The lookup of names in template instantiations does not conform to the rules in the standard in the following respects:

- Although only names from the template definition context are considered for names that are not functions, the lookup is not limited to those names visible at the point where the template is defined.
- Functions from the context where the template is referenced are considered for all function calls in the template. Functions from the referencing context are only visible for dependent function calls.

Argument-dependent lookup

When argument-dependent lookup is enabled, functions that are made visible using argument-dependent lookup can overload with those made visible by normal lookup. The standard requires that this overloading occur even when the name found by normal lookup is a block `extern` declaration. The compiler does this overloading, but in default mode, argument-dependent lookup is suppressed when the normal lookup finds a block `extern`.

This means a program can have different behavior, depending on whether it is compiled with or without argument-dependent lookup, even if the program makes no use of namespaces. For example:

```
struct A { };
A operator+(A, double);
void f()
{
    A a1;
    A operator+(A, int);
    a1 + 1.0;           // calls operator+(A, double) with arg-dependent lookup
}                       // enabled but otherwise calls operator+(A, int);
```

6.2.7 C++ exception handling

The ARM compilation tools fully support C++ exception handling. However, the compiler does not support this by default. You must enable C++ exception handling with the `--exceptions` option. See *--exceptions, --no_exceptions* on page 3-37 for more information.

Note

The Rogue Wave Standard C++ Library is provided with C++ exceptions enabled.

You can exercise limited control over exception table generation.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Function unwinding at runtime

By default, functions compiled with `--exceptions` can be unwound at runtime. See *--exceptions, --no_exceptions* on page 3-37 for more information. *Function unwinding* includes destroying C++ automatic variables, and restoring register values saved in the stack frame. Function unwinding is implemented by emitting an exception table describing the operations to be performed.

You can enable or disable unwinding for specific functions with the pragmas `#pragma exceptions_unwind` and `#pragma no_exceptions_unwind`, see *Pragmas* on page 5-47 for more information. The `--exceptions_unwind` option sets the initial value of this pragma.

Disabling function unwinding for a function has the following effects:

- Exceptions cannot be thrown through that function at runtime, and no stack unwinding occurs for that throw. If the throwing language is C++, then `std::terminate` is called.
- A very compact exception table representation can be used to describe the function, that assists smart linkers with table optimization.
- Function inlining is restricted, because the caller and callee must interact correctly.

Therefore, `#pragma no_exceptions_unwind` can be used to forcibly prevent unwinding in a way that requires no additional source decoration.

By contrast, in C++ an empty function exception specification permits unwinding as far as the protected function, then calls `std::unexpected()` in accordance with the ISO C++ Standard.

6.2.8 Extern inline functions

The ISO C++ Standard requires inline functions to be defined wherever you use them. To prevent the clashing of multiple out-of-line copies of inline functions, the C++ compiler emits out-of-line `extern` functions in common sections.

Out-of-line inline functions

The compiler emits inline functions out-of-line, in the following cases:

- The address of the function is taken, for example:

  ```
  inline int g()
  {
      return 1;
  }
  int (*fp)() = &g;
  ```

- The function cannot be inlined, for example, a recursive function:

  ```
  inline unsigned int fact(unsigned int n) {
      return n < 2 ? 1 : n * fact(n - 1);
  }
  ```

- The heuristic used by the compiler decides that it is better not to inline the function. This heuristic is influenced by `-Ospace` and `-Otime`. If you use `-Otime`, the compiler inlines more functions. You can override this heuristic by declaring a function with `__forceinline`. For example:

  ```
  __forceinline int g()
  {
      return 1;
  }
  ```

  See also *--forceinline* on page 3-40 for more information.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Appendix A
ARMv6 SIMD Instruction Intrinsics

This appendix describes the ARMv6 SIMD instruction intrinsics. It contains the following sections:

- *ARMv6 SIMD intrinsics by prefix* on page A-3
- *ARMv6 SIMD intrinsics, summary descriptions, byte lanes, affected flags* on page A-5
- *ARMv6 SIMD intrinsics, compatible processors and architectures* on page A-9
- *ARMv6 SIMD instruction intrinsics and APSR GE flags* on page A-10
- *__qadd16 intrinsic* on page A-11
- *__qadd8 intrinsic* on page A-12
- *__qasx intrinsic* on page A-13
- *__qsax intrinsic* on page A-14
- *__qsub16 intrinsic* on page A-15
- *__qsub8 intrinsic* on page A-16
- *__sadd16 intrinsic* on page A-17
- *__sadd8 intrinsic* on page A-18
- *__sasx intrinsic* on page A-19
- *__sel intrinsic* on page A-20
- *__shadd16 intrinsic* on page A-21
- *__shadd8 intrinsic* on page A-22
- *__shasx intrinsic* on page A-23
- *__shsax intrinsic* on page A-24
- *__shsub16 intrinsic* on page A-25
- *__shsub8 intrinsic* on page A-26
- *__smlad intrinsic* on page A-27

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

- *__smladx intrinsic* on page A-28
- *__smlald intrinsic* on page A-29
- *__smlaldx intrinsic* on page A-30
- *__smlsd intrinsic* on page A-31
- *__smlsdx intrinsic* on page A-32
- *__smlsld intrinsic* on page A-33
- *__smlsldx intrinsic* on page A-34
- *__smuad intrinsic* on page A-35
- *__smusd intrinsic* on page A-36
- *__smuadx intrinsic* on page A-38
- *__ssat16 intrinsic* on page A-39
- *__ssax intrinsic* on page A-40
- *__ssub16 intrinsic* on page A-41
- *__ssub8 intrinsic* on page A-42
- *__sxtab16 intrinsic* on page A-43
- *__sxtb16 intrinsic* on page A-44
- *__uadd16 intrinsic* on page A-45
- *__uadd8 intrinsic* on page A-46
- *__uasx intrinsic* on page A-47
- *__uhadd16 intrinsic* on page A-48
- *__uhadd8 intrinsic* on page A-49
- *__uhasx intrinsic* on page A-50
- *__uhsax intrinsic* on page A-51
- *__uhsub16 intrinsic* on page A-52
- *__uhsub8 intrinsic* on page A-53
- *__uqadd16 intrinsic* on page A-54
- *__uqadd8 intrinsic* on page A-55
- *__uqasx intrinsic* on page A-56
- *__uqsax intrinsic* on page A-57
- *__uqsub16 intrinsic* on page A-58
- *__uqsub8 intrinsic* on page A-59
- *__usad8 intrinsic* on page A-60
- *__usada8 intrinsic* on page A-61
- *__usax intrinsic* on page A-62
- *__usat16 intrinsic* on page A-63
- *__usub16 intrinsic* on page A-64
- *__usub8 intrinsic* on page A-65
- *__uxtab16 intrinsic* on page A-66
- *__uxtb16 intrinsic* on page A-67.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.1 ARMv6 SIMD intrinsics by prefix

Table A-1 shows the intrinsics according to prefix name.

The `__sel()` intrinsic falls outside the classifications shown in the table. This intrinsic selects bytes according to GE bit values.

Table A-1

	ARMv6 SIMD instruction intrinsics grouped by prefix					
Intrinsic classification	`__s`[a]	`__q`[b]	`__sh`[c]	`__u`[d]	`__uq`[e]	`__uh`[f]
Byte addition	`__sadd8`	`__qadd8`	`__shadd8`	`__uadd8`	`__uqadd8`	`__uhadd8`
Byte subtraction	`__ssub8`	`__qsub8`	`__shsub8`	`__usub8`	`__uqsub8`	`__uhsub8`
Halfword addition	`__sadd16`	`__qadd16`	`__shadd16`	`__uadd16`	`__uqadd16`	`__uhadd16`
Halfword subtraction	`__ssub16`	`__qsub16`	`__shsub16`	`__usub16`	`__uqsub16`	`__uhsub16`
Exchange halfwords within one operand, add high halfwords, subtract low halfwords	`__sasx`	`__qasx`	`__shasx`	`__uasx`	`__uqasx`	`__uhasx`
Exchange halfwords within one operand, subtract high halfwords, add low halfwords	`__ssax`	`__qsax`	`__shsax`	`__usax`	`__uqsax`	`__uhsax`
Unsigned sum of absolute difference	-	-	-	`__usad8`	-	-
Unsigned sum of absolute difference and accumulate	-	-	-	`__usada8`	-	-
Saturation to selected width	`__ssat16`	-	-	`__usat16`	-	-
Extract values (bit positions [23:16][7:0]), zero-extend to 16 bits	-	-	-	`__uxtb16`	-	-
Extract values (bit positions [23:16][7:0]) from second operand, zero-extend to 16 bits, add to first operand	-	-	-	`__uxtab16`	-	-
Sign-extend	`__sxtb16`	-	-	-	-	-
Sign-extend, add	`__sxtab16`	-	-	-	-	-
Signed multiply, add products	`__smuad`	-	-	-	-	-
Exchange halfwords of one operand, signed multiply, add products	`__smuadx`	-	-	-	-	-
Signed multiply, subtract products	`__smusd`	-	-	-	-	-
Exchange halfwords of one operand, signed multiply, subtract products	`__smusdx`	-	-	-	-	-
Signed multiply, add both results to another operand	`__smlad`	-	-	-	-	-
Exchange halfwords of one operand, perform 2x16-bit multiplication, add both results to another operand	`__smladx`	-	-	-	-	-
Perform 2x16-bit multiplication, add both results to another operand	`__smlald`	-	-	-	-	-

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table A-1 (continued)

Intrinsic classification	ARMv6 SIMD instruction intrinsics grouped by prefix					
	`__s`[a]	`__q`[b]	`__sh`[c]	`__u`[d]	`__uq`[e]	`__uh`[f]
Exchange halfwords of one operand, perform 2x16-bit multiplication, add both results to another operand	`__smlaldx`	-	-	-	-	-
Perform 2x16-bit signed multiplications, take difference of products, subtracting high halfword product from low halfword product, and add difference to a 32-bit accumulate operand	`__smlsd`	-	-	-	-	-
Exchange halfwords of one operand, perform two signed 16-bit multiplications, add difference of products to a 32-bit accumulate operand	`__smlsdx`	-	-	-	-	-
Perform 2x16-bit signed multiplications, take difference of products, subtracting high halfword product from low halfword product, add difference to a 64-bit accumulate operand	`__smlsld`	-	-	-	-	-
Exchange halfwords of one operand, perform 2x16-bit multiplications, add difference of products to a 64-bit accumulate operand	`__smlsldx`	-	-	-	-	-

a. Signed
b. Signed saturating
c. Signed halving
d. Unsigned
e. Unsigned saturating
f. Unsigned halving.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.2 ARMv6 SIMD intrinsics, summary descriptions, byte lanes, affected flags

Table A-2

Intrinsic	Summary description	Byte lanes		Affected flags
		Returns	Operands	
`__qadd16`	2 x 16-bit addition, saturated to range $-2^{15} \leq x \leq 2^{15} - 1$	int16x2	int16x2, int16x2	None
`__qadd8`	4 x 8-bit addition, saturated to range $-2^{7} \leq x \leq 2^{7} - 1$	int8x4	int8x4, int8x4	None
`__qasx`	Exchange halfwords of second operand, add high halfwords, subtract low halfwords, saturating in each case	int16x2	int16x2, int16x2	None
`__qsax`	Exchange halfwords of second operand, subtract high halfwords, add low halfwords, saturating in each case	int16x2	int16x2, int16x2	None
`__qsub16`	2 x 16-bit subtraction with saturation	int16x2	int16x2, int16x2	None
`__qsub8`	4 x 8-bit subtraction with saturation	int8x4	int8x4, int8x4	None
`__sadd16`	2 x 16-bit signed addition.	int16x2	int16x2, int16x2	APSR.GE bits
`__sadd8`	4 x 8-bit signed addition	int8x4	int8x4, int8x4	APSR.GE bits
`__sasx`	Exchange halfwords of second operand, add high halfwords, subtract low halfwords	int16x2	int16x2, int16x2	APSR.GE bits
`__sel`	Select each byte of the result from either the first operand or the second operand, according to the values of the GE bits. For each result byte, if the corresponding GE bit is set, the byte from the first operand is selected, otherwise the byte from the second operand is selected. Because of the way that int16x2 operations set two (duplicate) GE bits per value, the `__sel` intrinsic works equally well on (u)int16x2 and (u)int8x4 data.	uint8x4	uint8x4, uint8x4	None
`__shadd16`	2x16-bit signed addition, halving the results	int16x2	int16x2, int16x2	None
`__shadd8`	4x8-bit signed addition, halving the results	int8x4	int8x4, int8x4	None
`__shasx`	Exchange halfwords of the second operand, add high halfwords and subtract low halfwords, halving the results	int16x2	int16x2, int16x2	None
`__shsax`	Exchange halfwords of the second operand, subtract high halfwords and add low halfwords, halving the results	int16x2	int16x2, int16x2	None
`__shsub16`	2x16-bit signed subtraction, halving the results	int16x2	int16x2, int16x2	None
`__shsub8`	4x8-bit signed subtraction, halving the results	int8x4	int8x4, int8x4	None

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table A-2 (continued)

Intrinsic	Summary description	Byte lanes		Affected flags
		Returns	Operands	
`__smlad`	2x16-bit multiplication, adding both results to third operand	int32	int16x2, int16x2, int32	Q bit
`__smladx`	Exchange halfwords of the second operand, 2x16-bit multiplication, adding both results to third operand	int16x2	int16x2, int16x2	Q bit
`__smlald`	2x16-bit multiplication, adding both results to third operand. Overflow in addition is not detected.	int64	int16x2, int16x2, int64	None
`__smlaldx`	Exchange halfwords of second operand, perform 2x16-bit multiplication, adding both results to third operand. Overflow in addition is not detected.	int64	int16x2, int16x2, int64	None
`__smlsd`	2x16-bit signed multiplications. Take difference of products, subtract high halfword product from low halfword product, add difference to third operand.	int32	int16x2, int16x2, int32	Q bit
`__smlsdx`	Exchange halfwords of second operand, then 2x16-bit signed multiplications. Product difference is added to a third accumulate operand.	int32	int16x2, int16x2, int32	Q bit
`__smlsld`	2x16-bit signed multiplications. Take difference of products, subtracting high halfword product from low halfword product, and add difference to third operand. Overflow in addition is not detected.	int64	int16x2, int16x2, int64	None
`__smlsldx`	Exchange halfwords of second operand, then 2x16-bit signed multiplications. Take difference of products, subtracting high halfword product from low halfword product, and add difference to third operand. Overflow in addition is not detected.	int64	int16x2, int16x2, u64	None
`__smuad`	2x16-bit signed multiplications, adding the products together.	int32	int16x2, int16x2	Q bit
`__smusd`	2x16-bit signed multiplications. Take difference of products, subtracting high halfword product from low halfword product.	int32	int16x2, int16x2	None
`__smusdx`	2x16-bit signed multiplications. Product of high halfword of first operand and low halfword of second operand is subtracted from product of low halfword of first operand and high halfword of second operand, and difference is added to third operand.	int32	int16x2, int16x2	None
`__ssat16`	2x16-bit signed saturation to a selected width	int16x2	int16x2, /*constant*/ unsigned int	Q bit
`__ssax`	Exchange halfwords of second operand, subtract high halfwords and add low halfwords	int16x2	int16x2, int16x2	APSR.GE bits

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table A-2 (continued)

Intrinsic	Summary description	Byte lanes		Affected flags
		Returns	Operands	
`__ssub16`	2x16-bit signed subtraction	int16x2	int16x2, int16x2	APSR.GE bits
`__ssub8`	4x8-bit signed subtraction	int8x4	int8x4	APSR.GE bits
`__smuadx`	Exchange halfwords of second operand, perform 2x16-bit signed multiplications, and add products together	int32	int16x2, int16x2	Q bit
`__sxtab16`	Two values at bit positions [23:16][7:0] are extracted from second operand, sign-extended to 16 bits, and added to first operand	int16x2	int8x4, int16x2	None
`__sxtb16`	Two values at bit positions [23:16][7:0] are extracted from the operand and sign-extended to 16 bits	int16x2	int8x4	None
`__uadd16`	2x16-bit unsigned addition	uint16x2	uint16x2, uint16x2	APSR.GE bits
`__uadd8`	4x8-bit unsigned addition	uint8x4	uint8x4, uint8x4	APSR.GE bits
`__uasx`	Exchange halfwords of second operand, add high halfwords and subtract low halfwords	uint16x2	uint16x2, uint16x2	APSR.GE bits
`__uhadd16`	2x16-bit unsigned addition, halving the results	uint16x2	uint16x2, uint16x2	None
`__uhadd8`	4x8-bit unsigned addition, halving the results	uint8x4	uint8x4, uint8x4	None
`__uhasx`	Exchange halfwords of second operand, add high halfwords and subtract low halfwords, halving the results	uint16x2	uint16x2, uint16x2	None
`__uhsax`	Exchange halfwords of second operand, subtract high halfwords and add low halfwords, halving the results	uint16x2	uint16x2, uint16x2	None
`__uhsub16`	2x16-bit unsigned subtraction, halving the results	uint16x2	uint16x2, uint16x2	None
`__uhsub8`	4x8-bit unsigned subtraction, halving the results	uint8x4	uint8x4	None
`__uqadd16`	2x16-bit unsigned addition, saturating to range $0 \leq x \leq 2^{16} - 1$	uint16x2	uint16x2, uint16x2	None
`__uqadd8`	4x8-bit unsigned addition, saturating to range $0 \leq x \leq 2^{8} - 1$	uint8x4	uint8x4, uint8x4	None
`__uqasx`	Exchange halfwords of second operand, perform saturating unsigned addition on high halfwords and saturating unsigned subtraction on low halfwords	uint16x2	uint16x2, uint16x2	None
`__uqsax`	Exchange halfwords of second operand, perform saturating unsigned subtraction on high halfwords and saturating unsigned addition on low halfwords	uint16x2	uint16x2, uint16x2	None
`__uqsub16`	2x16-bit unsigned subtraction, saturating to range $0 \leq x \leq 2^{16} - 1$	uint16x2	uint16x2, uint16x2	None

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table A-2 (continued)

Intrinsic	Summary description	Byte lanes		Affected flags
		Returns	Operands	
`__uqsub8`	4x8-bit unsigned subtraction, saturating to range $0 \le x \le 2^8 - 1$	uint8x4	uint8x4, uint8x4	None
`__usad8`	4x8-bit unsigned subtraction, add absolute values of the differences together, return result as single unsigned integer	uint32	uint8x4, uint8x4	None
`__usada8`	4x8-bit unsigned subtraction, add absolute values of the differences together, and add result to third operand	uint32	uint8x4, uint8x4, uint32	None
`__usax`	Exchange halfwords of second operand, subtract high halfwords and add low halfwords	uint16x2	uint16x2, uint16x2	APSR.GE bits
`__usat16`	Saturate two 16-bit values to a selected unsigned range. Input values are signed and output values are non-negative.	int16x2	int16x2, /*constant*/ unsigned int	Q flag
`__usub16`	2x16-bit unsigned subtraction	uint16x2	uint16x2, uint16x2	APSR.GE bits
`__usub8`	4x8-bit unsigned subtraction	uint8x4	uint8x4, uint8x4	APSR.GE bits
`__uxtab16`	Two values at bit positions [23:16][7:0] are extracted from the second operand, zero-extended to 16 bits, and added to the first operand	uint16x2	uint8x4, uint16x2	None
`__uxtb16`	Two values at bit positions [23:16][7:0] are extracted from the operand and zero-extended to 16 bits	uint16x2	uint8x4	None

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.3 ARMv6 SIMD intrinsics, compatible processors and architectures

Table A-3 lists some ARMv6 SIMD instruction intrinsics and compatible processors and architectures, as examples of compatibility.

Use of intrinsics that are not available on your target platform results in linkage failure with undefined symbols.

Table A-3

Intrinsics	Compatible `--cpu` options
`__qadd16`, `__qadd8`, `__qasx`	`6, 6K, 6T2, 6Z, 7-A, 7-R, 7-A.security, Cortex-R4, Cortex-R4F, Cortex-R7, Cortex-R7.no_vfp, Cortex-M4, Cortex-M4.fp, MPCore, MPCore.no_vfp, MPCoreNoVFP, 88FR111, 88FR111.no_hw_divide, QSP, QSP.no_neon, QSP.no_neon.no_vfp`

A.3.1 See also

Reference

- *--cpu=list* on page 3-20
- *--cpu=name* on page 3-20.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.4 ARMv6 SIMD instruction intrinsics and APSR GE flags

Table A-4

Intrinsic	APSR.GE flag action	APSR.GE operation

`__sel`	Reads GE flags	`if APSR.GE[0] == 1 then res[7:0] = val1[7:0] else val2[7:0]`
		`if APSR.GE[1] == 1 then res[15:8] = val1[15:8] else val2[15:8]`
		`if APSR.GE[2] == 1 then res[23:16] = val1[23:16] else val2[23:16]`
		`if APSR.GE[3] == 1 then res[31:24] = val1[31:24] else val2[31:24]`

`__sadd16`	Sets or clears GE flags	`if sum1 • 0 then APSR.GE[1:0] = 11 else 00` `if sum2 • 0 then APSR.GE[3:2] = 11 else 00`

`__sadd8`	Sets or clears GE flags	`if sum1 • 0 then APSR.GE[0] = 1 else 0`
		`if sum2 • 0 then APSR.GE[1] = 1 else 0`
		`if sum3 • 0 then APSR.GE[2] = 1 else 0`
		`if sum4 • 0 then APSR.GE[3] = 1 else 0`

`__sasx`	Sets or clears GE flags	`if diff • 0 then APSR.GE[1:0] = 11 else 00` `if sum • 0 then APSR.GE[3:2] = 11 else 00`
`__ssax`	Sets or clears GE flags	`if sum • 0 then APSR.GE[1:0] = 11 else 00` `if diff • 0 then APSR.GE[3:2] = 11 else 00`
`__ssub16`	Sets or clears GE flags	`if diff1 • 0 then APSR.GE[1:0] = 11 else 00` `if diff2 • 0 then APSR.GE[3:2] = 11 else 00`

`__ssub8`	Sets or clears GE flags	`if diff1 • 0 then APSR.GE[0] = 1 else 0`
		`if diff2 • 0 then APSR.GE[1] = 1 else 0`
		`if diff3 • 0 then APSR.GE[2] = 1 else 0`
		`if diff4 • 0 then APSR.GE[3] = 1 else 0`

`__uadd16`	Sets or clears GE flags	`if sum1 • 0x10000 then APSR.GE[1:0] = 11 else 00` `if sum2 • 0x10000 then APSR.GE[3:2] = 11 else 00`

`__uadd8`	Sets or clears GE flags	`if sum1 • 0x100 then APSR.GE[0] = 1 else 0`
		`if sum2 • 0x100 then APSR.GE[1] = 1 else 0`
		`if sum3 • 0x100 then APSR.GE[2] = 1 else 0`
		`if sum4 • 0x100 then APSR.GE[3] = 1 else 0`

`__uasx`	Sets or clears GE flags	`if diff • 0 then APSR.GE[1:0] = 11 else 00` `if sum • 0x10000 then APSR.GE[3:2] = 11 else 00`
`__usax`	Sets or clears GE flags	`if sum • 0x10000 then APSR.GE[1:0] = 11 else 00` `if diff • 0 then APSR.GE[3:2] = 11 else 00`
`__usub16`	Sets or clears GE flags	`if diff1 • 0 then APSR.GE[1:0] = 11 else 00` `if diff2 • 0 then APSR.GE[3:2] = 11 else 00`

`__usub8`	Sets or clears GE flags	`if diff1 • 0 then APSR.GE[0] = 1 else 0`
		`if diff2 • 0 then APSR.GE[1] = 1 else 0`
		`if diff3 • 0 then APSR.GE[2] = 1 else 0`
		`if diff4 • 0 then APSR.GE[3] = 1 else 0`

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.5 __qadd16 intrinsic

This intrinsic inserts a QADD16 instruction into the instruction stream generated by the compiler. It enables you to perform two 16-bit integer arithmetic additions in parallel, saturating the results to the 16-bit signed integer range $-2^{15} \leq x \leq 2^{15} - 1$.

```
unsigned int __qadd16(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first two 16-bit summands

val2 holds the second two 16-bit summands.

The __qadd16 intrinsic returns:

- the saturated addition of the low halfwords in the low halfword of the return value
- the saturated addition of the high halfwords in the high halfword of the return value.

The returned results are saturated to the 16-bit signed integer range $-2^{15} \leq x \leq 2^{15} - 1$.

Example:

```
unsigned int add_halfwords(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __qadd16(val1, val2); /* res[15:0] = val1[15:0] + val2[15:0]
                                   res[16:31] = val1[31:16] + val2[31:16]
                                 */
    return res;
}
```

A.5.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Saturating instructions* on page 3-96 in the *Assembler Reference*
- *QADD, QSUB, QDADD, and QDSUB* on page 3-97 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.6 `__qadd8` intrinsic

This intrinsic inserts a `QADD8` instruction into the instruction stream generated by the compiler. It enables you to perform four 8-bit integer additions, saturating the results to the 8-bit signed integer range $-2^7 \leq x \leq 2^7 - 1$.

```
unsigned int __qadd8(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first four 8-bit summands

val2 holds the other four 8-bit summands.

The `__qadd8` intrinsic returns:

- the saturated addition of the first byte of each operand in the first byte of the return value
- the saturated addition of the second byte of each operand in the second byte of the return value
- the saturated addition of the third byte of each operand in the third byte of the return value
- the saturated addition of the fourth byte of each operand in the fourth byte of the return value.

The returned results are saturated to the 8-bit signed integer range $-2^7 \leq x \leq 2^7 - 1$.

Example:

```
unsigned int add_bytes(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __qadd8(val1,val2); /* res[7:0] = val1[7:0] + val2[7:0]
                                 res[15:8] = val1[15:8] + val2[15:8]
                                 res[23:16] = val1[23:16] + val2[23:16]
                                 res[31:24] = val1[31:24] + val2[31:24]
                              */
    return res;
}
```

A.6.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Saturating instructions* on page 3-96 in the *Assembler Reference*
- *QADD, QSUB, QDADD, and QDSUB* on page 3-97 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.7 __qasx intrinsic

This intrinsic inserts a QASX instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of the one operand, then add the high halfwords and subtract the low halfwords, saturating the results to the 16-bit signed integer range $-2^{15} \leq x \leq 2^{15} - 1$.

```
unsigned int __qasx(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first operand for the subtraction in the low halfword, and the first operand for the addition in the high halfword

val2 holds the second operand for the subtraction in the high halfword, and the second operand for the addition in the low halfword.

The __qasx intrinsic returns:

- the saturated subtraction of the high halfword in the second operand from the low halfword in the first operand, in the low halfword of the return value
- the saturated addition of the high halfword in the first operand and the low halfword in the second operand, in the high halfword of the return value.

The returned results are saturated to the 16-bit signed integer range $-2^{15} \leq x \leq 2^{15} - 1$.

Example:

```
unsigned int exchange_add_and_subtract(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __qasx(val1,val2); /* res[15:0] = val1[15:0] - val2[31:16]
                                res[31:16] = val1[31:16] + val2[15:0]
                              */
                             /* Alternative equivalent representation:
                                val2[15:0][31:16] = val2[31:16][15:0]
                                res[15:0] = val1[15:0] - val2[15:0]
                                res[31:16] = val[31:16] + val2[31:16]
                              */
    return res;
}
```

A.7.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Saturating instructions* on page 3-96 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.8 __qsax intrinsic

This intrinsic inserts a QSAX instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of one operand, then subtract the high halfwords and add the low halfwords, saturating the results to the 16-bit signed integer range $-2^{15} \leq x \leq 2^{15} - 1$.

```
unsigned int __qsax(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first operand for the addition in the low halfword, and the first operand for the subtraction in the high halfword

val2 holds the second operand for the addition in the high halfword, and the second operand for the subtraction in the low halfword.

The __qsax intrinsic returns:

- the saturated addition of the low halfword of the first operand and the high halfword of the second operand, in the low halfword of the return value
- the saturated subtraction of the low halfword of the second operand from the high halfword of the first operand, in the high halfword of the return value.

The returned results are saturated to the 16-bit signed integer range $-2^{15} \leq x \leq 2^{15} - 1$.

Example:

```
unsigned int exchange_subtract_and_add(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __qsax(val1,val2); /* res[15:0] = val1[15:0] + val2[31:16]
                                res[31:16] = val1[31:16] - val2[15:0]
                              */
                             /* Alternative equivalent representation:
                                val2[15:0][31:16] = val2[31:16][15:0]
                                res[15:0] = val1[15:0] + val2[15:0]
                                res[31:16] = val[31:16] - val2[31:16]
                              */
    return res;
}
```

A.8.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Saturating instructions* on page 3-96 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.9 `__qsub16` intrinsic

This intrinsic inserts a `QSUB16` instruction into the instruction stream generated by the compiler. It enables you to perform two 16-bit integer subtractions, saturating the results to the 16-bit signed integer range $-2^{15} \leq x \leq 2^{15} - 1$.

```
unsigned int __qsub16(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first halfword operands

val2 holds the second halfword operands.

The `__qsub16` intrinsic returns:

- the saturated subtraction of the low halfword in the second operand from the low halfword in the first operand, in the low halfword of the returned result
- the saturated subtraction of the high halfword in the second operand from the high halfword in the first operand, in the high halfword of the returned result.

The returned results are saturated to the 16-bit signed integer range $-2^{15} \leq x \leq 2^{15} - 1$.

Example:

```
unsigned int subtract_halfwords(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __qsub16(val1,val2); /* res[15:0] = val1[15:0] - val2[15:0]
                                  res[31:16] = val1[31:16] - val2[31:16]
                                */
    return res;
}
```

A.9.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Saturating instructions* on page 3-96 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.10 `__qsub8` intrinsic

This intrinsic inserts a `QSUB8` instruction into the instruction stream generated by the compiler. It enables you to perform four 8-bit integer subtractions, saturating the results to the 8-bit signed integer range $-2^7 \leq x \leq 2^7 - 1$.

```
unsigned int __qsub8(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first four 8-bit operands

val2 holds the second four 8-bit operands.

The `__qsub8` intrinsic returns:

- the subtraction of the first byte in the second operand from the first byte in the first operand, in the first byte of the return value
- the subtraction of the second byte in the second operand from the second byte in the first operand, in the second byte of the return value
- the subtraction of the third byte in the second operand from the third byte in the first operand, in the third byte of the return value
- the subtraction of the fourth byte in the second operand from the fourth byte in the first operand, in the fourth byte of the return value.

The returned results are saturated to the 8-bit signed integer range $-2^7 \leq x \leq 2^7 - 1$.

Example:

```
unsigned int subtract_bytes(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __qsub8(val1,val2); /* res[7:0] = val1[7:0] - val2[7:0]
                                 res[15:8] = val1[15:8] - val2[15:8]
                                 res[23:16] = val1[23:16] - val2[23:16]
                                 res[31:24] = val1[31:24] - val2[31:24]
                               */
    return res;
}
```

A.10.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Saturating instructions* on page 3-96 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.11 __sadd16 intrinsic

This intrinsic inserts an SADD16 instruction into the instruction stream generated by the compiler. It enables you to perform two 16-bit signed integer additions. The GE bits in the APSR are set according to the results of the additions.

```
unsigned int __sadd16(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first two 16-bit summands

val2 holds the second two 16-bit summands.

The __sadd16 intrinsic returns:

- the addition of the low halfwords in the low halfword of the return value
- the addition of the high halfwords in the high halfword of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *res* is the return value, then:

- if *res*[15:0] ≥ 0 then APSR.GE[1:0] = 11 else 00
- if *res*[31:16] ≥ 0 then APSR.GE[3:2] = 11 else 00.

Example:

```
unsigned int add_halfwords(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __sadd16(val1,val2); /* res[15:0] = val1[15:0] + val2[15:0]
                                  res[31:16] = val1[31:16] + val2[31:16]
                                */
    return res;
}
```

A.11.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *__sel intrinsic* on page A-20
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Saturating instructions* on page 3-96 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.12 `__sadd8` intrinsic

This intrinsic inserts an `SADD8` instruction into the instruction stream generated by the compiler. It enables you to perform four 8-bit signed integer additions. The GE bits in the APSR are set according to the results of the additions.

`unsigned int __sadd8(unsigned int val1, unsigned int val2)`

Where:

val1 holds the first four 8-bit summands

val2 holds the second four 8-bit summands.

The `__sadd8` intrinsic returns:

- the addition of the first bytes from each operand, in the first byte of the return value
- the addition of the second bytes of each operand, in the second byte of the return value
- the addition of the third bytes of each operand, in the third byte of the return value
- the addition of the fourth bytes of each operand, in the fourth byte of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *res* is the return value, then:

- if *res*[7:0] ≥ 0 then APSR.GE[0] = 1 else 0
- if *res*[15:8] ≥ 0 then APSR.GE[1] = 1 else 0.
- if *res*[23:16] ≥ 0 then APSR.GE[2] = 1 else 0.
- if *res*[31:24] ≥ 0 then APSR.GE[3] = 1 else 0.

Example:

```
unsigned int add_bytes(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __sadd16(val1,val2); /* res[7:0] = val1[7:0] + val2[7:0]
                                  res[15:8] = val1[15:8] + val2[15:8]
                                  res[23:16] = val1[23:16] + val2[23:16]
                                  res[31:24] = val1[31:24] + val2[31:24]
                                */
    return res;
}
```

A.12.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *__sel intrinsic* on page A-20
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Saturating instructions* on page 3-96 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.13 `__sasx` intrinsic

This intrinsic inserts an `SASX` instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of the second operand, add the high halfwords and subtract the low halfwords. The GE bits in the APRS are set according to the results.

```
unsigned int __sasx(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first operand for the subtraction in the low halfword, and the first operand for the addition in the high halfword

val2 holds the second operand for the subtraction in the high halfword, and the second operand for the addition in the low halfword.

The `__sasx` intrinsic returns:

- the subtraction of the high halfword in the second operand from the low halfword in the first operand, in the low halfword of the return value
- the addition of the high halfword in the first operand and the low halfword in the second operand, in the high halfword of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *res* is the return value, then:

- if *res*[15:0] ≥ 0 then APSR.GE[1:0] = 11 else 00
- if *res*[31:16] ≥ 0 then APSR.GE[3:2] = 11 else 00.

Example:

```
unsigned int exchange_subtract_add(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __sasx(val1,val2); /* res[15:0] = val1[15:0] - val2[31:16]
                                res[31:16] = val1[31:16] + val2[15:0]
                              */
    return res;
}
```

A.13.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *__sel intrinsic* on page A-20
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.14 `__sel` intrinsic

This intrinsic inserts a `SEL` instruction into the instruction stream generated by the compiler. It enables you to select bytes from the input parameters, whereby the bytes that are selected depend upon the results of previous SIMD instruction intrinsics. The results of previous SIMD instruction intrinsics are represented by the *Greater than or Equal* flags in the *Application Program Status Register* (APSR).

The `__sel` intrinsic works equally well on both halfword and byte operand intrinsic results. This is because halfword operand operations set two (duplicate) GE bits per value. For example, the `__sasx` intrinsic.

```
unsigned int __sel(unsigned int val1, unsigned int val2)
```

Where:

`val1` holds four selectable bytes

`val2` holds four selectable bytes.

The `__sel` intrinsic selects bytes from the input parameters and returns them in the return value, *`res`*, according to the following criteria:

```
if APSR.GE[0] == 1 then res[7:0] = val1[7:0] else res[7:0] = val2[7:0]
if APSR.GE[1] == 1 then res[15:8] = val1[15:8] else res[15:8] = val2[15:8]
if APSR.GE[2] == 1 then res[23:16] = val1[23:16] else res[23:16] = val2[23:16]
if APSR.GE[3] == 1 then res[31;24] = val1[31:24] else res = val2[31:24]
```

Example:

```
unsigned int ge_filter(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __sel(val1,val2);
    return res;
}

unsigned int foo(unsigned int a, unsigned int b)
{
  int res;
  int filtered_res;

    res = __sasx(a,b);  /* This intrinsic sets the GE flags */
    filtered_res = ge_filter(res); /* Filter the results of the __sasx */
                                   /* intrinsic. Some results are filtered */
                                   /* out based on the GE flags. */
    return filtered_res;
}
```

A.14.1 See also

- *__sadd16 intrinsic* on page A-17
- *__sasx intrinsic* on page A-19
- *__ssax intrinsic* on page A-40
- *__ssub8 intrinsic* on page A-42
- *__ssub16 intrinsic* on page A-41
- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SEL* on page 3-67 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.15 `__shadd16` intrinsic

This intrinsic inserts a `SHADD16` instruction into the instruction stream generated by the compiler. It enables you to perform two signed 16-bit integer additions, halving the results.

`unsigned int __shadd16(unsigned int val1, unsigned int val2)`

Where:

val1 holds the first two 16-bit summands

val2 holds the second two 16-bit summands.

The `__shadd16` intrinsic returns:

- the halved addition of the low halfwords from each operand, in the low halfword of the return value
- the halved addition of the high halfwords from each operand, in the high halfword of the return value.

Example:

```
unsigned int add_and_halve(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __shadd16(val1,val2); /* res[15:0] = (val1[15:0] + val2[15:0]) >> 1
                                   res[31:16] = (val1[31:16] + val2[31:16]) >> 1
                                */
    return res;
}
```

A.15.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.16 `__shadd8` intrinsic

This intrinsic inserts a `SHADD8` instruction into the instruction stream generated by the compiler. It enables you to perform four signed 8-bit integer additions, halving the results.

```
unsigned int __shadd8(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first four 8-bit summands

val2 holds the second four 8-bit summands.

The `__shadd8` intrinsic returns:

- the halved addition of the first bytes from each operand, in the first byte of the return value
- the halved addition of the second bytes from each operand, in the second byte of the return value
- the halved addition of the third bytes from each operand, in the third byte of the return value
- the halved addition of the fourth bytes from each operand, in the fourth byte of the return value.

Example:

```
unsigned int add_and_halve(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __shadd8(val1,val2); /* res[7:0] = (val1[7:0] + val2[7:0]) >> 1
                                  res[15:8] = (val1[15:8] + val2[15:8]) >> 1
                                  res[23:16] = (val1[23:16] + val2[23:16]) >> 1
                                  res[31:24] = (val1[31:24] + val2[31:24]) >> 1
                               */
    return res;
}
```

A.16.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.17 `__shasx` intrinsic

This intrinsic inserts a `SHASX` instruction into the instruction stream generated by the compiler. It enables you to exchange the two halfwords of one operand, perform one signed 16-bit integer addition and one signed 16-bit subtraction, and halve the results.

```
unsigned int __shasx(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first halfword operands

val2 holds the second halfword operands.

The `__shasx` intrinsic returns:

- the halved subtraction of the high halfword in the second operand from the low halfword in the first operand, in the low halfword of the return value
- the halved subtraction of the low halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

Example:

```
unsigned int exchange_add_subract_halve(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __shasx(val1,val2); /* res[15:0] = (val1[15:0] - val2[31:16]) >> 1
                                 res[31:16] = (val1[31:16] - val2[15:0]) >> 1
                               */
    return res;
}
```

A.17.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.18 __shsax intrinsic

This intrinsic inserts a SHSAX instruction into the instruction stream generated by the compiler. It enables you to exchange the two halfwords of one operand, perform one signed 16-bit integer subtraction and one signed 16-bit addition, and halve the results.

```
unsigned int __shsax(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first halfword operands

val2 holds the second halfword operands.

The __shsax intrinsic returns:

- the halved addition of the low halfword in the first operand and the high halfword in the second operand, in the low halfword of the return value
- the halved subtraction of the low halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

Example:

```
unsigned int exchange_subract_add_halve(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __shsax(val1,val2); /* res[15:0] = (val1[15:0] + val2[31:16]) >> 1
                                 res[31:16] = (val1[31:16] - val2[15:0]) >> 1
                               */
    return res;
}
```

A.18.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.19 `__shsub16` intrinsic

This intrinsic inserts a `SHSUB16` instruction into the instruction stream generated by the compiler. It enables you to perform two signed 16-bit integer subtractions, halving the results.

```
unsigned int __shsub16(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first halfword operands

val2 holds the second halfword operands.

The `__shsub16` intrinsic returns:

- the halved subtraction of the low halfword in the second operand from the low halfword in the first operand, in the low halfword of the return value
- the halved subtraction of the high halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

Example:

```
unsigned int add_and_halve(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __shsub16(val1,val2); /* res[15:0] = (val1[15:0] - val2[15:0]) >> 1
                                   res[31:16] = (val1[31:16] - val2[31:16]) >> 1
                                 */
    return res;
}
```

A.19.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.20 `__shsub8` intrinsic

This intrinsic inserts a `SHSUB8` instruction into the instruction stream generated by the compiler. It enables you to perform four signed 8-bit integer subtractions, halving the results.

```
unsigned int __shsub8(unsigned int val1, unsigned int val2)
```

Where:

`val1` holds the first four operands

`val2` holds the second four operands.

The `__shsub8` intrinsic returns:

- the halved subtraction of the first byte in the second operand from the first byte in the first operand, in the first byte of the return value
- the halved subtraction of the second byte in the second operand from the second byte in the first operand, in the second byte of the return value
- the halved subtraction of the third byte in the second operand from the third byte in the first operand, in the third byte of the return value
- the halved subtraction of the fourth byte in the second operand from the fourth byte in the first operand, in the fourth byte of the return value

Example:

```
unsigned int subtract_and_halve(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __shsub8(val1,val2); /* res[7:0] = (val1[7:0] - val2[7:0]) >> 1
                                  res[15:8] = (val1[15:8] - val2[15:8]) >> 1
                                  res[23:16] = (val1[23:16] - val2[23:16] >> 1
                                  res[31:24] = (val1[31:24] - val2[31:24] >> 1
                                */
    return res;
}
```

A.20.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.21 `__smlad` intrinsic

This intrinsic inserts an `SMLAD` instruction into the instruction stream generated by the compiler. It enables you to perform two signed 16-bit multiplications, adding both results to a 32-bit accumulate operand. The Q bit is set if the addition overflows. Overflow cannot occur during the multiplications.

```
unsigned int __smlad(unsigned int val1, unsigned int val2, unsigned int val3)
```

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication

val3 holds the accumulate value.

The `__smlad` intrinsic returns the product of each multiplication added to the accumulate value, as a 32-bit integer.

Example:

```
unsigned int dual_multiply_accumulate(unsigned int val1, unsigned int val2, unsigned
int val3)
{
  unsigned int res;

    res = __smlad(val1,val2,val3); /* p1 = val1[15:0] × val2[15:0]
                                      p2 = val1[31:16] × val2[31:16]
                                      res[31:0] = p1 + p2 + val3[31:0]
                                    */
    return res;
}
```

A.21.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMLAD and SMLSD* on page 3-89 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.22 __smladx intrinsic

This intrinsic inserts an SMLADX instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of the second operand, perform two signed 16-bit multiplications, adding both results to a 32-bit accumulate operand. The Q bit is set if the addition overflows. Overflow cannot occur during the multiplications.

```
unsigned int __smladx(unsigned int val1, unsigned int val2, unsigned int val3)
```

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication

val3 holds the accumulate value.

The __smladx intrinsic returns the product of each multiplication added to the accumulate value, as a 32-bit integer.

Example:

```
unsigned int dual_multiply_accumulate(unsigned int val1, unsigned int val2, unsigned
int val3)
{
  unsigned int res;

    res = __smladx(val1,val2,val3); /* p1 = val1[15:0] × val2[31:16]
                                       p2 = val1[31:16] × val2[15:0]
                                       res[31:0] = p1 + p2 + val3[31:0]
                                     */
    return res;
}
```

A.22.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMLAD and SMLSD* on page 3-89 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.23 `__smlald` intrinsic

This intrinsic inserts an `SMLALD` instruction into the instruction stream generated by the compiler. It enables you to perform two signed 16-bit multiplications, adding both results to a 64-bit accumulate operand. Overflow is only possible as a result of the 64-bit addition. This overflow is not detected if it occurs. Instead, the result wraps around modulo 2^{64}.

```
unsigned long long__smlald(unsigned int val1, unsigned int val2, unsigned long long
val3)
```

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication

val3 holds the accumulate value.

The `__smlald` intrinsic returns the product of each multiplication added to the accumulate value.

Example:

```
unsigned int dual_multiply_accumulate(unsigned int val1, unsigned int val2, unsigned
int val3)
{
  unsigned int res;

    res = __smlald(val1,val2,val3); /* p1 = val1[15:0] × val2[15:0]
                                       p2 = val1[31:16] × val2[31:16]
                                       sum = p1 + p2 + val3[63:32][31:0]
                                       res[63:32] = sum[63:32]
                                       res[31:0] = sum[31:0]
                                     */
    return res;
}
```

A.23.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMLALD and SMLSLD* on page 3-91 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.24 `__smlaldx` intrinsic

This intrinsic inserts an `SMLALDX` instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of the second operand, and perform two signed 16-bit multiplications, adding both results to a 64-bit accumulate operand. Overflow is only possible as a result of the 64-bit addition. This overflow is not detected if it occurs. Instead, the result wraps around modulo 2^{64}.

```
unsigned long long__smlaldx(unsigned int val1, unsigned int val2, unsigned long long
val3)
```

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication

val3 holds the accumulate value.

The `__smlald` intrinsic returns the product of each multiplication added to the accumulate value.

Example:

```
unsigned int dual_multiply_accumulate(unsigned int val1, unsigned int val2, unsigned
int val3)
{
  unsigned int res;

    res = __smlald(val1,val2,val3); /* p1 = val1[15:0] × val2[31:16]
                                       p2 = val1[31:16] × val2[15:0]
                                       sum = p1 + p2 + val3[63:32][31:0]
                                       res[63:32] = sum[63:32]
                                       res[31:0] = sum[31:0]
                                     */
    return res;
}
```

A.24.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMLALD and SMLSLD* on page 3-91 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.25 `__smlsd` intrinsic

This intrinsic inserts an `SMLSD` instruction into the instruction stream generated by the compiler. It enables you to perform two 16-bit signed multiplications, take the difference of the products, subtracting the high halfword product from the low halfword product, and add the difference to a 32-bit accumulate operand. The Q bit is set if the accumulation overflows. Overflow cannot occur during the multiplications or the subtraction.

```
unsigned int__smlsd(unsigned int val1, unsigned int val2, unsigned int val3)
```

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication

val3 holds the accumulate value.

The `__smlsd` intrinsic returns the difference of the product of each multiplication, added to the accumulate value.

Example:

```
unsigned int dual_multiply_diff_prods(unsigned int val1, unsigned int val2, unsigned
int val3)
{
  unsigned int res;

    res = __smlsd(val1,val2,val3); /* p1 = val1[15:0] × val2[15:0]
                                      p2 = val1[31:16] × val2[31:16]
                                      res[31:0] = p1 - p2 + val3[31:0]
                                    */
    return res;
}
```

A.25.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMLAD and SMLSD* on page 3-89 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.26 `__smlsdx` intrinsic

This intrinsic inserts an `SMLSDX` instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords in the second operand, then perform two 16-bit signed multiplications. The difference of the products is added to a 32-bit accumulate operand. The Q bit is set if the addition overflows. Overflow cannot occur during the multiplications or the subtraction.

`unsigned int__smlsdx(unsigned int val1, unsigned int val2, unsigned int val3)`

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication

val3 holds the accumulate value.

The `__smlsd` intrinsic returns the difference of the product of each multiplication, added to the accumulate value.

Example:

```
unsigned int dual_multiply_diff_prods(unsigned int val1, unsigned int val2, unsigned
int val3)
{
  unsigned int res;

    res = __smlsd(val1,val2,val3); /* p1 = val1[15:0] × val2[31:16]
                                      p2 = val1[31:16] × val2[15:0]
                                      res[31:0] = p1 - p2 + val3[31:0]
                                    */
    return res;
}
```

A.26.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMLAD and SMLSD* on page 3-89 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.27 `__smlsld` intrinsic

This intrinsic inserts an `SMLSLD` instruction into the instruction stream generated by the compiler. It enables you to perform two 16-bit signed multiplications, take the difference of the products, subtracting the high halfword product from the low halfword product, and add the difference to a 64-bit accumulate operand. Overflow cannot occur during the multiplications or the subtraction. Overflow can occur as a result of the 64-bit addition, and this overflow is not detected. Instead, the result wraps round to modulo 2^{64}.

```
unsigned long long__smlsld(unsigned int val1, unsigned int val2, unsigned long long val3)
```

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication

val3 holds the accumulate value.

The `__smlsld` intrinsic returns the difference of the product of each multiplication, added to the accumulate value.

Example:

```
unsigned long long dual_multiply_diff_prods(unsigned int val1, unsigned int val2,
unsigned long long val3)
{
  unsigned int res;

    res = __smlsld(val1,val2,val3); /* p1 = val1[15:0] × val2[15:0]
                                       p2 = val1[31:16] × val2[31:16]
                                       res[63:0] = p1 - p2 + val3[63:0]
                                     */
    return res;
}
```

A.27.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMLALD and SMLSLD* on page 3-91 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.28 `__smlsldx` intrinsic

This intrinsic inserts an `SMLSLDX` instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of the second operand, perform two 16-bit multiplications, adding the difference of the products to a 64-bit accumulate operand. Overflow cannot occur during the multiplications or the subtraction. Overflow can occur as a result of the 64-bit addition, and this overflow is not detected. Instead, the result wraps round to modulo 2^{64}.

```
unsigned long long__smlsldx(unsigned int val1, unsigned int val2, unsigned long long
val3)
```

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication

val3 holds the accumulate value.

The `__smlsld` intrinsic returns the difference of the product of each multiplication, added to the accumulate value.

Example:

```
unsigned long long dual_multiply_diff_prods(unsigned int val1, unsigned int val2,
unsigned long long val3)
{
  unsigned int res;

    res = __smlsld(val1,val2,val3); /* p1 = val1[15:0] × val2[31:16]
                                       p2 = val1[31:16] × val2[15:0]
                                       res[63:0] = p1 - p2 + val3[63:0]
                                     */
    return res;
}
```

A.28.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMLALD and SMLSLD* on page 3-91 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.29 `__smuad` intrinsic

This intrinsic inserts an `SMUAD` instruction into the instruction stream generated by the compiler. It enables you to perform two 16-bit signed multiplications, adding the products together. The Q bit is set if the addition overflows.

```
unsigned int__smuad(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication.

The `__smuad` intrinsic returns the products of the two 16-bit signed multiplications.

Example:

```
unsigned int dual_multiply_prods(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __smuad(val1,val2); /* p1 = val1[15:0] × val2[15:0]
                                 p2 = val1[31:16] × val2[31:16]
                                 res[31:0] = p1 + p2
                               */
    return res;
}
```

A.29.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMUAD{X} and SMUSD{X}* on page 3-85 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.30 `__smusd` intrinsic

This intrinsic inserts an `SMUSD` instruction into the instruction stream generated by the compiler. It enables you to perform two 16-bit signed multiplications, taking the difference of the products by subtracting the high halfword product from the low halfword product.

```
unsigned int__smusd(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication.

The `__smusd` intrinsic returns the difference of the products of the two 16-bit signed multiplications.

Example:

```
unsigned int dual_multiply_prods(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __smuad(val1,val2); /* p1 = val1[15:0] × val2[15:0]
                                 p2 = val1[31:16] × val2[31:16]
                                 res[31:0] = p1 - p2
                               */
    return res;
}
```

A.30.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMUAD{X} and SMUSD{X}* on page 3-85 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.31 `__smusdx` intrinsic

This intrinsic inserts an `SMUSDX` instruction into the instruction stream generated by the compiler. It enables you to perform two 16-bit signed multiplications, subtracting one of the products from the other. The halfwords of the second operand are exchanged before performing the arithmetic. This produces top × bottom and bottom × top multiplication.

```
unsigned int__smusdx(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication.

The `__smusdx` intrinsic returns the difference of the products of the two 16-bit signed multiplications.

Example:

```
unsigned int dual_multiply_prods(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __smuad(val1,val2); /* p1 = val1[15:0] × val2[31:16]
                                 p2 = val1[31:16] × val2[15:0]
                                 res[31:0] = p1 - p2
                               */
    return res;
}
```

A.31.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMUAD{X} and SMUSD{X}* on page 3-85 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.32 __smuadx intrinsic

This intrinsic inserts an SMUADX instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of the second operand, perform two 16-bit signed integer multiplications, and add the products together. Exchanging the halfwords of the second operand produces top × bottom and bottom × top multiplication. The Q flag is set if the addition overflows. The multiplications cannot overflow.

```
unsigned int__smuadx(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first halfword operands for each multiplication

val2 holds the second halfword operands for each multiplication.

The __smuadx intrinsic returns the products of the two 16-bit signed multiplications.

Example:

```
unsigned int exchange_dual_multiply_prods(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __smuadx(val1,val2); /* val2[31:16][15:0] = val2[15:0][31:16]
                                  p1 = val1[15:0] × val2[15:0]
                                  p2 = val1[31:16] × val2[31:16]
                                  res[31:0] = p1 + p2
                                */
    return res;
}
```

A.32.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SMUAD{X} and SMUSD{X}* on page 3-85 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.33 __ssat16 intrinsic

This intrinsic inserts an SSAT16 instruction into the instruction stream generated by the compiler. It enables you to saturate two signed 16-bit values to a selected signed range.

The Q bit is set if either operation saturates.

`unsigned int __saturate_halfwords(unsigned int val1, unsigned int val2)`

Where:

val1 holds the two signed 16-bit values to be saturated

val2 is the bit position for saturation, an integral constant expression in the range 1 to 16.

The __ssat16 intrinsic returns:

- the signed saturation of the low halfword in *val1*, saturated to the bit position specified in *val2* and returned in the low halfword of the return value
- the signed saturation of the high halfword in *val1*, saturated to the bit position specified in *val2* and returned in the high halfword of the return value.

Example:

```
unsigned int saturate_halfwords(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __ssat16(val1,val2); /* Saturate halfwords in val1 to the signed
                                  range specified by the bit position in val2 */
    return res;
}
```

A.33.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Saturating instructions* on page 3-96 in the *Assembler Reference*
- *SSAT16 and USAT16* on page 3-106 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.34 `__ssax` intrinsic

This intrinsic inserts an `SSAX` instruction into the instruction stream generated by the compiler. It enables you to exchange the two halfwords of one operand and perform one 16-bit integer subtraction and one 16-bit addition.

The GE bits in the APSR are set according to the results.

`unsigned int __ssax(unsigned int val1, unsigned int val2)`

Where:

val1 holds the first operand for the addition in the low halfword, and the first operand for the subtraction in the high halfword

val2 holds the second operand for the addition in the high halfword, and the second operand for the subtraction in the low halfword.

The `__ssax` intrinsic returns:

- the addition of the low halfword in the first operand and the high halfword in the second operand, in the low halfword of the return value
- the subtraction of the low halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *res* is the return value, then:

- if *res*[15:0] ≥ 0 then APSR.GE[1:0] = 11 else 00
- if *res*[31:16] ≥ 0 then APSR.GE[3:2] = 11 else 00.

Example:

```
unsigned int exchange_subtract_add(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __ssax(val1,val2); /* res[15:0] = val1[15:0] + val2[31:16]
                                res[31:16] = val1[31:16] - val2[15:0]
                              */
    return res;
}
```

A.34.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.35 `__ssub16` intrinsic

This intrinsic inserts an `SSUB16` instruction into the instruction stream generated by the compiler. It enables you to perform two 16-bit signed integer subtractions.

The GE bits in the APSR are set according to the results.

`unsigned int __ssub16(unsigned int val1, unsigned int val2)`

Where:

val1 holds the first operands of each subtraction in the low and the high halfwords

val2 holds the second operands for each subtraction in the low and the high halfwords.

The `__ssub16` intrinsic returns:

- the subtraction of the low halfword in the second operand from the low halfword in the first operand, in the low halfword of the return value
- the subtraction of the high halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *res* is the return value, then:

- if *res*[15:0] ≥ 0 then APSR.GE[1:0] = 11 else 00
- if *res*[31:16] ≥ 0 then APSR.GE[3:2] = 11 else 00.

Example:

```
unsigned int subtract halfwords(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __ssub16(val1,val2); /* res[15:0] = val1[15:0] - val2[15:0]
                                  res[31:16] = val1[31:16] - val2[31:16]
                                */
    return res;
}
```

A.35.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *__sel intrinsic* on page A-20
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.36 `__ssub8` intrinsic

This intrinsic inserts an `SSUB8` instruction into the instruction stream generated by the compiler. It enables you to perform four 8-bit signed integer subtractions.

The GE bits in the APSR are set according to the results.

`unsigned int __ssub8(unsigned int val1, unsigned int val2)`

Where:

val1 holds the first four 8-bit operands of each subtraction

val2 holds the second four 8-bit operands of each subtraction.

The `__ssub8` intrinsic returns:

- the subtraction of the first byte in the second operand from the first byte in the first operand, in the first bytes of the return value
- the subtraction of the second byte in the second operand from the second byte in the first operand, in the second byte of the return value
- the subtraction of the third byte in the second operand from the third byte in the first operand, in the third byte of the return value
- the subtraction of the fourth byte in the second operand from the fourth byte in the first operand, in the fourth byte of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *res* is the return value, then:

- if *res*[8:0] ≥ 0 then APSR.GE[0] = 1 else 0
- if *res*[15:8] ≥ 0 then APSR.GE[1] = 1 else 0
- if *res*[23:16] ≥ 0 then APSR.GE[2] = 1 else 0
- if *res*[31:24] ≥ 0 then APSR.GE[3] = 1 else 0.

Example:

```
unsigned int subtract bytes(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __ssub8(val1,val2); /* res[7:0] = val1[7:0] - val2[7:0]
                                 res[15:8] = val1[15:8] - val2[15:8]
                                 res[23:16] = val1[23:16] - val2[23:16]
                                 res[31:24] = val1[31:24] - val2[31:24]
                               */
    return res;
}
```

A.36.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *__sel intrinsic* on page A-20
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.37 `__sxtab16` intrinsic

This intrinsic inserts an `SXTAB16` instruction into the instruction stream generated by the compiler. It enables you to extract two 8-bit values from the second operand (at bit positions [7:0] and [23:16]), sign-extend them to 16-bits each, and add the results to the first operand.

`unsigned int __sxtab16(unsigned int val1, unsigned int val2)`

Where:

val1 holds the values that the extracted and sign-extended values are added to

val2 holds the two 8-bit values to be extracted and sign-extended.

The `__sxtab16` intrinsic returns the addition of *val1* and *val2*, where the 8-bit values in *val2*[7:0] and *val2*[23:16] have been extracted and sign-extended prior to the addition.

Example:

```
unsigned int extract_sign_extend_and_add(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __sxtab16(val1,val2); /* res[15:0]
                                    = val1[15:0] + SignExtended(val2[7:0])

                                   res[31:16]
                                    = val1[31:16] + SignExtended(val2[23:16])
                                 */
    return res;
}
```

A.37.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SXT, SXTA, UXT, and UXTA* on page 3-111 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.38 `__sxtb16` intrinsic

This intrinsic inserts an `SXTB16` instruction into the instruction stream generated by the compiler. It enables you to extract two 8-bit values from an operand and sign-extend them to 16 bits each.

```
unsigned int __sxtb16(unsigned int val)
```

Where *val*[7:0] and *val*[23:16] hold the two 8-bit values to be sign-extended.

The `__sxtb16` intrinsic returns the 8-bit values sign-extended to 16-bit values.

Example:

```
unsigned int sign_extend(unsigned int val)
{
  unsigned int res;

    res = __sxtb16(val1,val2); /* res[15:0] = SignExtended(val[7:0]
                                  res[31:16] = SignExtended(val[23:16]
                                */
    return res;
}
```

A.38.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SXT, SXTA, UXT, and UXTA* on page 3-111 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.39 __uadd16 intrinsic

This intrinsic inserts a UADD16 instruction into the instruction stream generated by the compiler. It enables you to perform two 16-bit unsigned integer additions.

The GE bits in the APSR are set according to the results.

```
unsigned int __uadd16(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first two halfword summands for each addition

val2 holds the second two halfword summands for each addition.

The __uadd16 intrinsic returns:

- the addition of the low halfwords in each operand, in the low halfword of the return value
- the addition of the high halfwords in each operand, in the high halfword of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *res* is the return value, then:

- if *res*[15:0] ≥ 0x10000 then APSR.GE[0] = 11 else 00
- if *res*[31:16] ≥ 0x10000 then APSR.GE[1] = 11 else 00.

Example:

```
unsigned int add_halfwords(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uadd16(val1,val2); /* res[15:0] = val1[15:0] + val2[15:0]
                                  res[31:16] = val1[31:16] + val2[31:16]
                                */
    return res;
}
```

A.39.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.40 `__uadd8` intrinsic

This intrinsic inserts a `UADD8` instruction into the instruction stream generated by the compiler. It enables you to perform four unsigned 8-bit integer additions.

The GE bits in the APSR are set according to the results.

`unsigned int __uadd8(unsigned int val1, unsigned int val2)`

Where:

`val1` holds the first four 8-bit summands for each addition

`val2` holds the second four 8-bit summands for each addition.

The `__uadd8` intrinsic returns:

- the addition of the first bytes in each operand, in the first byte of the return value
- the addition of the second bytes in each operand, in the second byte of the return value
- the addition of the third bytes in each operand, in the third byte of the return value
- the addition of the fourth bytes in each operand, in the fourth byte of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *`res`* is the return value, then:

- if *res*[7:0] ≥ 0x100 then APSR.GE[0] = 1 else 0
- if *res*[15:8] ≥ 0x100 then APSR.GE[1] = 1 else 0
- if *res*[23:16] ≥ 0x100 then APSR.GE[2] = 1 else 0
- if *res*[31:24] ≥ 0x100 then APSR.GE[3] = 1 else 0.

Example:

```
unsigned int add_bytes(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uadd8(val1,val2); /* res[7:0] = val1[7:0] + val2[7:0]
                                 res[15:8] = val1[15:8] + val2[15:8]
                                 res[23:16] = val1[23:16] + val2[23:16]
                                 res[31:24] = val1[31:24] + val2[31:24]
                               */
    return res;
}
```

A.40.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.41 `__uasx` intrinsic

This intrinsic inserts a `UASX` instruction into the instruction stream generated by the compiler. It enables you to exchange the two halfwords of the second operand, add the high halfwords and subtract the low halfwords.

The GE bits in the APSR are set according to the results.

`unsigned int __uasx(unsigned int val1, unsigned int val2)`

Where:

val1 holds the first operand for the subtraction in the low halfword, and the first operand for the addition in the high halfword

val2 holds the second operand for the subtraction in the high halfword and the second operand for the addition in the low halfword.

The `__uasx` intrinsic returns:

- the subtraction of the high halfword in the second operand from the low halfword in the first operand, in the low halfword of the return value
- the addition of the high halfword in the first operand and the low halfword in the second operand, in the high halfword of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *res* is the return value, then:

- if *res*[15:0] ≥ 0 then APSR.GE[1:0] = 11 else 00
- if *res*[31:16] ≥ 0x10000 then APSR.GE[3:2] = 11 else 00.

Example:

```
unsigned int exchange_add_subtract(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uasx(val1,val2); /* res[15:0] = val1[15:0] - val2[31:16]
                                res[31:16] = val1[31:16] + val2[15:0]
                              */
    return res;
}
```

A.41.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.42 `__uhadd16` intrinsic

This intrinsic inserts a `UHADD16` instruction into the instruction stream generated by the compiler. It enables you to perform two unsigned 16-bit integer additions, halving the results.

```
unsigned int __uhadd16(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first two 16-bit summands

val2 holds the second two 16-bit summands.

The `__uhadd16` intrinsic returns:

- the halved addition of the low halfwords in each operand, in the low halfword of the return value
- the halved addition of the high halfwords in each operand, in the high halfword of the return value.

Example:

```
unsigned int add_halfwords_then_halve(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uhadd16(val1,val2); /* res[15:0] = (val1[15:0] + val2[15:0]) >> 1
                                   res[31:16] = (val1[31:16] + val2[31:16]) >> 1
                                 */
    return res;
}
```

A.42.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.43 __uhadd8 intrinsic

This intrinsic inserts a UHADD8 instruction into the instruction stream generated by the compiler. It enables you to perform four unsigned 8-bit integer additions, halving the results.

```
unsigned int __uhadd8(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first four 8-bit summands

val2 holds the second four 8-bit summands.

The __uhadd8 intrinsic returns:

- the halved addition of the first bytes in each operand, in the first byte of the return value
- the halved addition of the second bytes in each operand, in the second byte of the return value
- the halved addition of the third bytes in each operand, in the third byte of the return value
- the halved addition of the fourth bytes in each operand, in the fourth byte of the return value.

Example:

```
unsigned int add_bytes_then_halve(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uhadd8(val1,val2); /* res[7:0] = (val1[7:0] + val2[7:0]) >> 1
                                  res[15:8] = (val1[15:8] + val2[15:8]) >> 1
                                  res[23:16] = (val1[23:16] + val2[23:16]) >> 1
                                  res[31:24] = (val1[31:24] + val2[31:24]) >> 1
                                */
    return res;
}
```

A.43.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.44 `__uhasx` intrinsic

This intrinsic inserts a `UHASX` instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of the second operand, add the high halfwords and subtract the low halfwords, halving the results.

```
unsigned int __uhasx(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first operand for the subtraction in the low halfword, and the first operand for the addition in the high halfword

val2 holds the second operand for the subtraction in the high halfword, and the second operand for the addition in the low halfword.

The `__uhasx` intrinsic returns:

- the halved subtraction of the high halfword in the second operand from the low halfword in the first operand
- the halved addition of the high halfword in the first operand and the low halfword in the second operand.

Example:

```
unsigned int exchange_add_subtract(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uhasx(val1,val2); /* res[15:0] = (val1[15:0] - val2[31:16]) >> 1
                                 res[31:16] = (val1[31:16] + val2[15:0]) >> 1
                                */
    return res;
}
```

A.44.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.45 `__uhsax` intrinsic

This intrinsic inserts a `UHSAX` instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of the second operand, subtract the high halfwords and add the low halfwords, halving the results.

```
unsigned int __uhsax(unsigned int val1, unsigned int val2)
```

Where:

`val1` holds the first operand for the addition in the low halfword, and the first operand for the subtraction in the high halfword

`val2` holds the second operand for the addition in the high halfword, and the second operand for the subtraction in the low halfword.

The `__uhsax` intrinsic returns:

- the halved addition of the high halfword in the second operand and the low halfword in the first operand, in the low halfword of the return value
- the halved subtraction of the low halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

Example:

```
unsigned int exchange_subtract_add(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uhsax(val1,val2); /* res[15:0] = (val1[15:0] + val2[31:16]) >> 1
                                 res[31:16] = (val1[31:16] - val2[15:0]) >> 1
                               */
    return res;
}
```

A.45.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.46 `__uhsub16` intrinsic

This intrinsic inserts a `UHSUB16` instruction into the instruction stream generated by the compiler. It enables you to perform two unsigned 16-bit integer subtractions, halving the results.

```
unsigned int __uhsub16(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first two 16-bit operands

val2 holds the second two 16-bit operands.

The `__uhsub16` intrinsic returns:

- the halved subtraction of the low halfword in the second operand from the low halfword in the first operand, in the low halfword of the return value
- the halved subtraction of the high halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

Example:

```
unsigned int subtract_and_halve(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uhsub16(val1,val2); /* res[15:0] = (val1[15:0] + val2[15:0]) >> 1
                                   res[31:16] = (val1[31:16] - val2[31:16]) >> 1
                                 */
    return res;
}
```

A.46.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.47 `__uhsub8` intrinsic

This intrinsic inserts a `UHSUB8` instruction into the instruction stream generated by the compiler. It enables you to perform four unsigned 8-bit integer subtractions, halving the results.

`unsigned int __uhsub8(unsigned int val1, unsigned int val2)`

Where:

val1 holds the first four 8-bit operands

val2 holds the second four 8-bit operands.

The `__uhsub8` intrinsic returns:

- the halved subtraction of the first byte in the second operand from the first byte in the first operand, in the first byte of the return value
- the halved subtraction of the second byte in the second operand from the second byte in the first operand, in the second byte of the return value
- the halved subtraction of the third byte in the second operand from the third byte in the first operand, in the third byte of the return value
- the halved subtraction of the fourth byte in the second operand from the fourth byte in the first operand, in the fourth byte of the return value.

Example:

```
unsigned int subtract_and_halve(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uhsub8(val1,val2); /* res[7:0] = (val1[7:0] - val2[7:0]) >> 1
                                  res[15:8] = (val1[15:8] - val2[15:8]) >> 1
                                  res[23:16] = (val1[23:16] - val2[23:16]) >> 1
                                  res[31:24] = (val1[31:24] - val2[31:24]) >> 1
                                */
    return res;
}
```

A.47.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.48 `__uqadd16` intrinsic

This intrinsic inserts a `UQADD16` instruction into the instruction stream generated by the compiler. It enables you to perform two unsigned 16-bit integer additions, saturating the results to the 16-bit unsigned integer range $0 \leq x \leq 2^{16} - 1$.

```
unsigned int __uqadd16(unsigned int val1, unsigned int val2)
```

Where:

`val1` holds the first two halfword summands

`val2` holds the second two halfword summands.

The `__uqadd16` intrinsic returns:

- the addition of the low halfword in the first operand and the low halfword in the second operand
- the addition of the high halfword in the first operand and the high halfword in the second operand, in the high halfword of the return value.

The results are saturated to the 16-bit unsigned integer range $0 \leq x \leq 2^{16} - 1$.

Example:

```
unsigned int add_halfwords(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uqadd16(val1,val2); /* res[15:0] = val1[15:0] + val2[15:0]
                                   res[31:16] = val1[31:16] + val2[31:16]
                                 */
    return res;
}
```

A.48.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.49 `__uqadd8` intrinsic

This intrinsic inserts a `UQADD8` instruction into the instruction stream generated by the compiler. It enables you to perform four unsigned 8-bit integer additions, saturating the results to the 8-bit unsigned integer range $0 \leq x \leq 2^8 - 1$.

```
unsigned int __uqadd8(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first four 8-bit summands

val2 holds the second four 8-bit summands.

The `__uqadd8` intrinsic returns:

- the addition of the first bytes in each operand, in the first byte of the return value
- the addition of the second bytes in each operand, in the second byte of the return value
- the addition of the third bytes in each operand, in the third byte of the return value
- the addition of the fourth bytes in each operand, in the fourth byte of the return value.

The results are saturated to the 8-bit unsigned integer range $0 \leq x \leq 2^8 - 1$.

Example:

```
unsigned int add_bytes(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uqadd8(val1,val2); /* res[7:0] = val1[7:0] + val2[7:0]
                                  res[15:8] = val1[15:8] + val2[15:8]
                                  res[23:16] = val1[23:16] + val2[23:16]
                                  res[31:24] = val1[31:24] + val2[31:24]
                                */
    return res;
}
```

A.49.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.50 `__uqasx` intrinsic

This intrinsic inserts a `UQASX` instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of the second operand and perform one unsigned 16-bit integer addition and one unsigned 16-bit subtraction, saturating the results to the 16-bit unsigned integer range $0 \leq x \leq 2^{16} - 1$.

`unsigned int __uqasx(unsigned int val1, unsigned int val2)`

Where:

val1 holds the first two halfword operands

val2 holds the second two halfword operands.

The `__uqasx` intrinsic returns:

- the subtraction of the high halfword in the second operand from the low halfword in the first operand, in the low halfword of the return value
- the subtraction of the low halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

The results are saturated to the 16-bit unsigned integer range $0 \leq x \leq 2^{16} - 1$.

Example:

```
unsigned int exchange_add_subtract(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uqasx(val1,val2); /* res[15:0] = val1[15:0] - val2[31:16]
                                 res[31:16] = val1[31:16] + val2[15:0]
                               */
    return res;
}
```

A.50.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.51 `__uqsax` intrinsic

This intrinsic inserts a `UQSAX` instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of the second operand and perform one unsigned 16-bit integer subtraction and one unsigned 16-bit addition, saturating the results to the 16-bit unsigned integer range $0 \leq x \leq 2^{16} - 1$.

```
unsigned int __uqsax(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first 16-bit operand for the addition in the low halfword, and the first 16-bit operand for the subtraction in the high halfword

val2 holds the second 16-bit halfword for the addition in the high halfword, and the second 16-bit halfword for the subtraction in the low halfword.

The `__uqsax` intrinsic returns:

- the addition of the low halfword in the first operand and the high halfword in the second operand, in the low halfword of the return value
- the subtraction of the low halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

The results are saturated to the 16-bit unsigned integer range $0 \leq x \leq 2^{16} - 1$.

Example:

```
unsigned int exchange_subtract_add(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uqsax(val1,val2); /* res[15:0] = val1[15:0] + val2[31:16]
                                 res[31:16] = val1[31:16] - val2[15:0]
                              */
    return res;
}
```

A.51.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.52 `__uqsub16` intrinsic

This intrinsic inserts a `UQSUB16` instruction into the instruction stream generated by the compiler. It enables you to perform two unsigned 16-bit integer subtractions, saturating the results to the 16-bit unsigned integer range $0 \leq x \leq 2^{16} - 1$.

```
unsigned int __uqsub16(unsigned int val1, unsigned int val2)
```

Where:

`val1` holds the first halfword operands for each subtraction

`val2` holds the second halfword operands for each subtraction.

The `__uqsub16` intrinsic returns:

- the subtraction of the low halfword in the second operand from the low halfword in the first operand, in the low halfword of the return value
- the subtraction of the high halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

The results are saturated to the 16-bit unsigned integer range $0 \leq x \leq 2^{16} - 1$.

Example:

```
unsigned int subtract_halfwords(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uqsub16(val1,val2); /* res[15:0] = val1[15:0] - val2[15:0]
                                   res[31:16] = val1[31:16] - val2[31:16]
                                */
    return res;
}
```

A.52.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.53 `__uqsub8` intrinsic

This intrinsic inserts a `UQSUB8` instruction into the instruction stream generated by the compiler. It enables you to perform four unsigned 8-bit integer subtractions, saturating the results to the 8-bit unsigned integer range $0 \leq x \leq 2^8 - 1$.

```
unsigned int __uqsub8(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first four 8-bit operands

val2 holds the second four 8-bit operands.

The `__uqsub8` intrinsic returns:

- the subtraction of the first byte in the second operand from the first byte in the first operand, in the first byte of the return value
- the subtraction of the second byte in the second operand from the second byte in the first operand, in the second byte of the return value
- the subtraction of the third byte in the second operand from the third byte in the first operand, in the third byte of the return value
- the subtraction of the fourth byte in the second operand from the fourth byte in the first operand, in the fourth byte of the return value.

The results are saturated to the 8-bit unsigned integer range $0 \leq x \leq 2^8 - 1$.

Example:

```
unsigned int subtract_bytes(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uqsub8(val1,val2); /* res[7:0] = val1[7:0] - val2[7:0]
                                  res[15:8] = val1[15:8] - val2[15:8]
                                  res[23:16] = val1[23:16] - val2[23:16]
                                  res[31:24] = val1[31:24] - val2[31:24]
                               */
    return res;
}
```

A.53.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.54 `__usad8` intrinsic

This intrinsic inserts a `USAD8` instruction into the instruction stream generated by the compiler. It enables you to perform four unsigned 8-bit subtractions, and add the absolute values of the differences together, returning the result as a single unsigned integer.

```
unsigned int __usad8(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first four 8-bit operands for the subtractions

val2 holds the second four 8-bit operands for the subtractions.

The `__usad8` intrinsic returns the sum of the absolute differences of:

- the subtraction of the first byte in the second operand from the first byte in the first operand
- the subtraction of the second byte in the second operand from the second byte in the first operand
- the subtraction of the third byte in the second operand from the third byte in the first operand
- the subtraction of the fourth byte in the second operand from the fourth byte in the first operand.

The sum is returned as a single unsigned integer.

Example:

```
unsigned int subtract_add_abs(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __usad8(val1,val2); /* absdiff1 = val1[7:0] - val2[7:0]
                                 absdiff2 = val1[15:8] - val2[15:8]
                                 absdiff3 = val1[23:16] - val2[23:16]
                                 absdiff4 = val1[31:24] - val2[31:24]
                                 res[31:0] = absdiff1 + absdiff2 + absdiff3
                                  + absdiff4
                              */
    return res;
}
```

A.54.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *USAD8 and USADA8* on page 3-104 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.55 __usada8 intrinsic

This intrinsic inserts a USADA8 instruction into the instruction stream generated by the compiler. It enables you to perform four unsigned 8-bit subtractions, and add the absolute values of the differences to a 32-bit accumulate operand.

```
unsigned int __usada8(unsigned int val1, unsigned int val2, unsigned int val3)
```

Where:

val1 holds the first four 8-bit operands for the subtractions

val2 holds the second four 8-bit operands for the subtractions

val3 holds the accumulation value.

The __usada8 intrinsic returns the sum of the absolute differences of the following bytes, added to the accumulation value:

- the subtraction of the first byte in the second operand from the first byte in the first operand
- the subtraction of the second byte in the second operand from the second byte in the first operand
- the subtraction of the third byte in the second operand from the third byte in the first operand
- the subtraction of the fourth byte in the second operand from the fourth byte in the first operand.

Example:

```
unsigned int subtract_add_diff_accumulate(unsigned int val1, unsigned int val2,
unsigned int val3)
{
  unsigned int res;

    res = __usada8(val1,val2,val3); /* absdiff1 = val1[7:0] - val2[7:0]
                                       absdiff2 = val1[15:8] - val2[15:8]
                                       absdiff3 = val1[23:16] - val2[23:16]
                                       absdiff4 = val1[31:24] - val2[31:24]
                                       sum = absdiff1 + absdiff2 + absdiff3
                                        + absdiff4
                                       res[31:0] = sum[31:0] + val3[31:0]
                              */
    return res;
}
```

A.55.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *USAD8 and USADA8* on page 3-104 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.56 `__usax` intrinsic

This intrinsic inserts a `USAX` instruction into the instruction stream generated by the compiler. It enables you to exchange the halfwords of the second operand, subtract the high halfwords and add the low halfwords.

The GE bits in the APSR are set according to the results.

`unsigned int __usax(unsigned int val1, unsigned int val2)`

Where:

val1 holds the first operand for the addition in the low halfword, and the first operand for the subtraction in the high halfword

val2 holds the second operand for the addition in the high halfword, and the second operand for the subtraction in the low halfword.

The `__usax` intrinsic returns:

- the addition of the low halfword in the first operand and the high halfword in the second operand, in the low halfword of the return value
- the subtraction of the low halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *res* is the return value, then:

- if *res*[15:0] $\geq$ 0x10000 then APSR.GE[1:0] = 11 else 00
- if *res*[31:16] $\geq$ 0 then APSR.GE[3:2] = 11 else 00.

Example:

```
unsigned int exchange_subtract_add(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __usax(val1,val2); /* res[15:0] = val1[15:0] + val2[31:16]
                                res[31:16] = val1[31:16] - val2[15:0]
                             */
    return res;
}
```

A.56.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.57 `__usat16` intrinsic

This intrinsic inserts a `USAT16` instruction into the instruction stream generated by the compiler. It enables you to saturate two signed 16-bit values to a selected unsigned range. The Q flag is set if either operation saturates.

```
unsigned int __usat16(unsigned int val1, /* constant */ unsigned int val2)
```

Where:

val1 holds the two 16-bit values that are to be saturated

val2 specifies the bit position for saturation, and must be an integral constant expression.

The `__usat16` intrinsic returns the saturation of the two signed 16-bit values, as non-negative values.

Example:

```
unsigned int saturate_halfwords(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __usax(val1,val2); /* Saturate halfwords in val1 to the unsigned
                                range specified by the bit position in val2
                              */    return res;
}
```

A.57.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SSAT16 and USAT16* on page 3-106 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.58 `__usub16` intrinsic

This intrinsic inserts a `USUB16` instruction into the instruction stream generated by the compiler. It enables you to perform two 16-bit unsigned integer subtractions.

The GE bits in the APSR are set according to the results.

```
unsigned int __usub16(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first two halfword operands

val2 holds the second two halfword operands.

The `__usub16` intrinsic returns:

- the subtraction of the low halfword in the second operand from the low halfword in the first operand, in the low halfword of the return value
- the subtraction of the high halfword in the second operand from the high halfword in the first operand, in the high halfword of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *res* is the return value, then:

- if *res*[15:0] ≥ 0 then APSR.GE[1:0] = 11 else 00
- if *res*[31:16] ≥ 0 then APSR.GE[3:2] = 11 else 00.

Example:

```
unsigned int subtract_halfwords(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __usub16(val1,val2); /* res[15:0] = val1[15:0] - val2[15:0]
                                  res[31:16] = val1[31:16] - val2[31:16]
                                */
}
```

A.58.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.59 __usub8 intrinsic

This intrinsic inserts a USUB8 instruction into the instruction stream generated by the compiler. It enables you to perform four 8-bit unsigned integer subtractions.

The GE bits in the APSR are set according to the results.

```
unsigned int __usub8(unsigned int val1, unsigned int val2)
```

Where:

val1 holds the first four 8-bit operands

val2 holds the second four 8-bit operands.

The __usub8 intrinsic returns:

- the subtraction of the first byte in the second operand from the first byte in the first operand, in the first byte of the return value
- the subtraction of the second byte in the second operand from the second byte in the first operand, in the second byte of the return value
- the subtraction of the third byte in the second operand from the third byte in the first operand, in the third byte of the return value
- the subtraction of the fourth byte in the second operand from the fourth byte in the first operand, in the fourth byte of the return value.

Each bit in APSR.GE is set or cleared for each byte in the return value, depending on the results of the operation. If *res* is the return value, then:

- if *res*[7:0] ≥ 0 then APSR.GE[0] = 1 else 0
- if *res*[15:8] ≥ 0 then APSR.GE[1] = 1 else 0
- if *res*[23:16] ≥ 0 then APSR.GE[2] = 1 else 0
- if *res*[31:24] ≥ 0 then APSR.GE[3] = 1 else 0.

Example:

```
unsigned int subtract(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __usub18(val1,val2); /* res[7:0] = val1[7:0] - val2[7:0]
                                  res[15:8] = val1[15:8] - val2[15:8]
                                  res[23:16] = val1[23:16] - val2[23:16]
                                  res[31:24] = val1[31:24] - val2[31:24]
                                */
}
```

A.59.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *Parallel add and subtract* on page 3-102 in the *Assembler Reference*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.60 `__uxtab16` intrinsic

This intrinsic inserts a `UXTAB16` instruction into the instruction stream generated by the compiler. It enables you to extract two 8-bit values from one operand, zero-extend them to 16 bits each, and add the results to two 16-bit values from another operand.

```
unsigned int __uxtab16(unsigned int val1, unsigned int val2)
```

Where *val2*[7:0] and *val2*[23:16] hold the two 8-bit values to be zero-extended.

The `__uxtab16` intrinsic returns the 8-bit values in *val2*, zero-extended to 16-bit values and added to *val1*.

Example:

```
unsigned int extend_add(unsigned int val1, unsigned int val2)
{
  unsigned int res;

    res = __uxtab16(val1,val2); /* res[15:0] = ZeroExt(val2[7:0] to 16 bits)
                                     + val1[15:0]
                                   res[31:16] = ZeroExt(val2[31:16] to 16 bits)
                                     + val1[31:16]
                                 */
    return res;
}
```

A.60.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SXT, SXTA, UXT, and UXTA* on page 3-111 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

A.61 __uxtb16 intrinsic

This intrinsic inserts a UXTB16 instruction into the instruction stream generated by the compiler. It enables you to extract two 8-bit values from an operand and zero-extend them to 16 bits each.

```
unsigned int __uxtb16(unsigned int val)
```

Where *val*[7:0] and *val*[23:16] hold the two 8-bit values to be zero-extended.

The __uxtb16 intrinsic returns the 8-bit values zero-extended to 16-bit values.

Example:

```
unsigned int zero_extend(unsigned int val)
{
  unsigned int res;

    res = __uxtb16(val1,val2); /* res[15:0] = ZeroExtended(val[7:0])
                                  res[31:16] = ZeroExtended(val[23:16])
                                */
    return res;
}
```

A.61.1 See also

- *ARMv6 SIMD intrinsics* on page 5-88
- *Instruction summary* on page 3-2 in the *Assembler Reference*
- *SXT, SXTA, UXT, and UXTA* on page 3-111 in the *Assembler Reference.*

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Appendix B
Via File Syntax

This appendix describes the syntax of via files accepted by all the ARM development tools. It contains the following sections:

- *Overview of via files* on page B-2
- *Syntax* on page B-3.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

B.1 Overview of via files

Via files are plain text files that contain command-line arguments and options to ARM development tools. You can use via files with all the ARM command-line tools, that is, you can specify a via file from the command line using the `--via` command-line option with:

- `armcc`
- `armasm`
- `armlink`
- `fromelf`
- `armar`.

See the documentation for the individual tool for more information.

Note

In general, you can use a via file to specify any command-line option to a tool, including `--via`. This means that you can call multiple nested via files from within a via file.

This section includes:

- *Via file evaluation*.

B.1.1 Via file evaluation

When a tool that supports via files is invoked it:

1. Replaces the first specified `--via` *`via_file`* argument with the sequence of argument words extracted from the via file, including recursively processing any nested `--via` commands in the via file.
2. Processes any subsequent `--via` *`via_file`* arguments in the same way, in the order they are presented.

That is, via files are processed in the order you specify them, and each via file is processed completely including processing nested via files before processing the next via file.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

B.2 Syntax

Via files must conform to the following syntax rules:

- A via file is a text file containing a sequence of words. Each word in the text file is converted into an argument string and passed to the tool.
- Words are separated by whitespace, or the end of a line, except in delimited strings. For example:

  `--c90 --strict` (two words)

  `--c90--strict` (one word)
- The end of a line is treated as whitespace. For example:

  ```
  --c90
  --strict
  ```

  is equivalent to:

  `--c90 --strict`
- Strings enclosed in quotation marks ("), or apostrophes (') are treated as a single word. Within a quoted word, an apostrophe is treated as an ordinary character. Within an apostrophe delimited word, a quotation mark is treated as an ordinary character.

  Quotation marks are used to delimit filenames or path names that contain spaces. For example:

  `-I C:\My Project\includes` (three words) `-I "C:\My Project\includes"` (two words)

  Apostrophes can be used to delimit words that contain quotes. For example:

  `-DNAME='"RealView Compilation Tools"'` (one word)
- Characters enclosed in parentheses are treated as a single word. For example:

  `--option(x, y, z)` (one word)

  `--option (x, y, z)` (two words)
- Within quoted or apostrophe delimited strings, you can use a backslash (\) character to escape the quote, apostrophe, and backslash characters.
- A word that occurs immediately next to a delimited word is treated as a single word. For example:

  `-I"C:\Project\includes"`

  is treated as the single word:

  `-IC:\Project\includes`
- Lines beginning with a semicolon (`;`) or a hash (`#`) character as the first nonwhitespace character are comment lines. If a semicolon or hash character appears anywhere else in a line, it is not treated as the start of a comment. For example:

  ```
  -o objectname.axf      ;this is not a comment
  ```

  A comment ends at the end of a line, or at the end of the file. There are no multi-line comments, and there are no part-line comments.
- Lines that include the preprocessor option `-Dsymbol="value"` must be delimited with a single quote, either as `'-Dsymbol="value"'` or as `-Dsymbol='"value"'`. For example:

  ```
  -c -DFOO_VALUE='"FOO_VALUE"'
  ```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Appendix C
Summary Table of GNU Language Extensions

GNU provides many extensions to the C and C++ languages. These extensions are also supported by the ARM compiler in GNU mode (for example GNU C90). Some extensions are supported in a non GNU mode (for example C90). This appendix lists the language extensions that the ARM compiler supports in non GNU modes and the modes they are supported in. The Origin column shows whether the language feature is part of any of the C90, C99, or C++ ISO Standards. The Origin column shows GCC-specific if the feature originated as a GCC extension.

Table C-1 Supported GNU extensions

GNU extension	Origin	Modes supported
__alignof__ on page 5-3	GCC-Specific.	C90, C99, C++.
Aggregate initializer elements for automatic variables	Standard C99, Standard C++.	C99, C++.
`asm` keyword	Standard C++.	C++.
Assembler labels	-	C90, C99, C++.
Compound literals	Standard C99.	C99.
Designated initializers	Standard C99.	C99.
Extended lvalues	Standard C++.	C++.
Function attributes on page 5-25	-	C90, C99, C++.
Inline functions	Standard C99, Standard C++.	C99, C++.
__attribute__((aligned)) variable attribute on page 5-41	GCC-specific.	C90, C99, C++.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table C-1 Supported GNU extensions (continued)

GNU extension	Origin	Modes supported
__attribute__((deprecated)) variable attribute on page 5-41	GCC-specific.	C90, C99, C++.
__attribute__((packed)) variable attribute on page 5-42	GCC-specific.	C90, C99.
__attribute__((section("name"))) variable attribute on page 5-42	GCC-specific.	C99.
__attribute__((unused)) variable attribute on page 5-43	GCC-specific.	C90, C99, C++.
__attribute__((used)) variable attribute on page 5-43	GCC-specific.	C90, C99.
__attribute__((weak)) variable attribute on page 5-45	GCC-specific.	C90, C99, C++.
Variadic macros	Standard C99.	C90, C99, C++.

Other information

- *Which GNU language extensions are supported by the ARM Compiler?*, http://infocenter.arm.com/help/index.jsp?topic=/com.arm.doc.faqs/ka14717.html

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Appendix D
Standard C Implementation Definition

This appendix gives information required by the ISO C standard for conforming C implementations. It contains the following section:

- *Implementation definition* on page D-2
- *Behaviors considered undefined by the ISO C Standard* on page D-8.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

D.1 Implementation definition

Appendix G of the ISO C standard (ISO/IEC 9899:1990 (E)) collates information about portability issues. Sub-clause G3 lists the behavior that each implementation must document.

Note

This appendix does not duplicate information that is part of Chapter 5 *Compiler-specific Features*. This appendix provides references where applicable.

The following subsections correspond to the relevant sections of sub-clause G3. They describe aspects of the ARM C compiler and C library, not defined by the ISO C standard, that are implementation-defined:

Note

The support for the `wctype.h` and `wchar.h` headers excludes wide file operations.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

D.1.1 Translation

Diagnostic messages produced by the compiler are of the form:

source-file, *line-number*: *severity*: *error-code*: *explanation*

where *severity* is one of:

`[blank]` If the severity is blank, this is a remark and indicates common, but sometimes unconventional, use of C or C++. Remarks are not displayed by default. Use the `--remarks` option to display remark messages. Compilation continues.

`Warning` Flags unusual conditions in your code that might indicate a problem. Compilation continues.

`Error` Indicates a problem that causes the compilation to stop. For example, violations in the syntactic or semantic rules of the C or C++ language.

`Internal fault`

Indicates an internal problem with the compiler. Contact your supplier with the information listed in Chapter 1 *Conventions and Feedback*.

Here:

error-code Is a number identifying the error type.

explanation Is a text description of the error.

See Chapter 6 *Compiler Diagnostic Messages* in *Using the Compiler* for more information.

D.1.2 Environment

The mapping of a command line from the ARM architecture-based environment into arguments to `main()` is implementation-specific. The generic ARM C library supports the following:

- *main()*
- *Interactive device* on page D-4
- *Redirecting standard input, output, and error streams* on page D-4.

main()

The arguments given to `main()` are the words of the command line not including input/output redirections, delimited by whitespace, except where the whitespace is contained in double quotes.

Note

- A whitespace character is any character where the result of `isspace()` is true.
- A double quote or backslash character \ inside double quotes must be preceded by a backslash character.
- An input/output redirection is not recognized inside double quotes.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Interactive device

In a nonhosted implementation of the ARM C library, the term *interactive device* might be meaningless. The generic ARM C library supports a pair of devices, both called `:tt`, intended to handle keyboard input and VDU screen output. In the generic implementation:

- no buffering is done on any stream connected to `:tt` unless input/output redirection has occurred
- if input/output redirection other than to `:tt` has occurred, full file buffering is used except that line buffering is used if both `stdout` and `stderr` were redirected to the same file.

Redirecting standard input, output, and error streams

Using the generic ARM C library, the standard input, output and error streams can be redirected at runtime. For example, if `mycopy` is a program running on a host debugger that copies the standard input to the standard output, the following line runs the program:

```
mycopy < infile > outfile 2> errfile
```

and redirects the files as follows:

`stdin` The standard input stream is redirected to `infile`.

`stdout` The standard output stream is redirected to `outfile`.

`stderr` The standard error stream is redirected to `errfile`.

The permitted redirections are:

`0< filename` Reads `stdin` from *`filename`*.

`< filename` Reads `stdin` from *`filename`*.

`1> filename` Writes `stdout` to *`filename`*.

`> filename` Writes `stdout` to *`filename`*.

`2> filename` Writes `stderr` to *`filename`*.

`2>&1` Writes `stderr` to the same place as `stdout`.

`>& file` Writes both `stdout` and `stderr` to *`filename`*.

`>> filename` Appends `stdout` to *`filename`*.

`>>& filename` Appends both `stdout` and `stderr` to *`filename`*.

To redirect `stdin`, `stdout`, and `stderr` on the target, you must define:

```
#pragma import(_main_redirection)
```

File redirection is done only if either:

- the invoking operating system supports it
- the program reads and writes characters and has not replaced the C library functions `fputc()` and `fgetc()`.

D.1.3 Identifiers

See *Character sets and identifiers* on page 6-2 for more information.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

D.1.4 Characters

See *Character sets and identifiers* on page 6-2 for more information.

D.1.5 Integers

See *Integer* on page 6-4 for more information.

D.1.6 Floating-point

See *Float* on page 6-4 for more information.

D.1.7 Arrays and pointers

See *Arrays and pointers* on page 6-4 for more information.

D.1.8 Registers

Using the ARM compiler, you can declare any number of local objects to have the storage class **`register`**.

D.1.9 Structures, unions, enumerations, and bitfields

The ISO/IEC C standard requires the following implementation details to be documented for structured data types:

- the outcome when a member of a union is accessed using a member of different type
- the padding and alignment of members of structures
- whether a plain **`int`** bitfield is treated as a **`signed int`** bitfield or as an **`unsigned int`** bitfield
- the order of allocation of bitfields within a unit
- whether a bitfield can straddle a storage-unit boundary
- the integer type chosen to represent the values of an enumeration type.

See Chapter 6 *C and C++ Implementation Details* for more information.

Unions

See *Unions* on page 6-6 for information.

Enumerations

See *Enumerations* on page 6-6 for information.

Padding and alignment of structures

See *Structures* on page 6-7 for information.

Bitfields

See *Bitfields* on page 6-8 for information.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

D.1.10 Qualifiers

An object that has a volatile-qualified type is accessed as a word, halfword, or byte as determined by its size and alignment. For volatile objects larger than a word, the order of accesses to the parts of the object is undefined. Updates to volatile bitfields generally require a read-modify-write. Accesses to aligned word, halfword and byte types are atomic. Other volatile accesses are not necessarily atomic.

Otherwise, reads and writes to volatile qualified objects occur as directly implied by the source code, in the order implied by the source code.

D.1.11 Expression evaluation

The compiler can re-order expressions involving only associative and commutative operators of equal precedence, even in the presence of parentheses. For example, `a + (b + c)` might be evaluated as `(a + b) + c` if `a`, `b`, and `c` are integer expressions.

Between sequence points, the compiler can evaluate expressions in any order, regardless of parentheses. Therefore, side effects of expressions between sequence points can occur in any order.

The compiler can evaluate function arguments in any order.

Any aspect of evaluation order not prescribed by the relevant standard can be varied by:

- the optimization level you are compiling at
- the release of the compiler you are using.

D.1.12 Preprocessing directives

The ISO standard C header files can be referred to as described in the standard, for example, `#include <stdio.h>`.

Quoted names for includable source files are supported. The compiler accepts host filenames or UNIX filenames. For UNIX filenames on non-UNIX hosts, the compiler tries to translate the filename to a local equivalent.

The recognized `#pragma` directives are shown in *Pragmas* on page 5-47.

D.1.13 Library functions

The ISO C library variants are listed in *C and C++ runtime libraries* on page 2-6 in *Using ARM® C and C++ Libraries and Floating-Point Support*.

The precise nature of each C library is unique to the particular implementation. The generic ARM C library has, or supports, the following features:

- The macro `NULL` expands to the integer constant 0.
- If a program redefines a reserved external identifier such as `printf`, an error might occur when the program is linked with the standard libraries. If it is not linked with standard libraries, no error is detected.
- The `__aeabi_assert()` function prints details of the failing diagnostic on `stderr` and then calls the `abort()` function:

  ```
  *** assertion failed: expression, file name, line number
  ```


Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Note

The behavior of the `assert` macro depends on the conditions in operation at the most recent occurrence of `#include <assert.h>`. See *Program exit and the assert macro* on page 2-62 in *Using ARM® C and C++ Libraries and Floating-Point Support* for more information.

For implementation details of mathematical functions, macros, locale, signals, and input/output see Chapter 2 *The ARM C and C++ libraries* in *Using ARM® C and C++ Libraries and Floating-Point Support*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

D.2 Behaviors considered undefined by the ISO C Standard

The following are considered undefined behavior by the ISO C Standard:

- In character and string escapes, if the character following the \ has no special meaning, the value of the escape is the character itself. For example, a warning is generated if you use `\s` because it is the same as `s`.
- A **`struct`** that has no named fields but at least one unnamed field is accepted by default, but generates an error in strict 1990 ISO Standard C.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Appendix E
Standard C++ Implementation Definition

The ARM compiler supports the majority of the language features described in the ISO/IEC standard for C++ when compiling C++. This appendix lists the C++ language features defined in the standard, and states whether or not that language feature is supported by ARM C++. It contains the following sections:

- *Integral conversion* on page E-2
- *Calling a pure virtual function* on page E-3
- *Major features of language support* on page E-4
- *Standard C++ library implementation definition* on page E-5.

Note

This appendix does not duplicate information that is part of the standard C implementation. See Appendix D *Standard C Implementation Definition*.

When compiling C++ in ISO C mode, the ARM compiler is identical to the ARM C compiler. Where there is an implementation feature specific to either C or C++, this is noted in the text. For extensions to Standard C++, see:

- *Standard C++ language extensions* on page 4-13
- *C99 language features available in C++ and C90* on page 4-6
- *Standard C and Standard C++ language extensions* on page 4-16.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

E.1 Integral conversion

During integral conversion, if the destination type is signed, the value is unchanged if it can be represented in the destination type and bitfield width. Otherwise, the value is truncated to fit the size of the destination type.

Note

This section is related to Section 4.7 Integral conversions, in the ISO/IEC standard.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

E.2 Calling a pure virtual function

Calling a pure virtual function is illegal. If your code calls a pure virtual function, then the compiler includes a call to the library function `__cxa_pure_virtual`.

`__cxa_pure_virtual` raises the signal **SIGPVFN**. The default signal handler prints an error message and exits. See *__default_signal_handler()* on page 2-8 in the *ARM® C and C++ Libraries and Floating-Point Support Reference* for more information.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

E.3 Major features of language support

Table E-1 shows the major features of the language supported by this release of ARM C++.

Table E-1 Major feature support for language

Major feature	ISO/IEC standard section	Support
Core language	1 to 13	Yes.
Templates	14	Yes, with the exception of export templates.
Exceptions	15	Yes.
Libraries	17 to 27	See the *Standard C++ library implementation definition* on page E-5, the *ARM® C and C++ Libraries and Floating-Point Support Reference*, and *Using ARM® C and C++ Libraries and Floating-Point Support*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

E.4 Standard C++ library implementation definition

Version 2.02.03 of the Rogue Wave library provides a subset of the library defined in the standard. There are small differences from the 1999 ISO C standard. For information on the implementation definition, see *Standard C++ library implementation definition* on page 2-114 in *Using ARM® C and C++ Libraries and Floating-Point Support*.

The library can be used with user-defined functions to produce target-dependent applications. See *C and C++ runtime libraries* on page 2-6 in *Using ARM® C and C++ Libraries and Floating-Point Support*.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Appendix F
C and C++ Compiler Implementation Limits

This appendix lists the implementation limits when using the ARM compiler to compile C and C++. It contains the following sections:

- *C++ ISO/IEC standard limits* on page F-2
- *Limits for integral numbers* on page F-4
- *Limits for floating-point numbers* on page F-5.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

F.1 C++ ISO/IEC standard limits

The ISO/IEC C++ standard recommends minimum limits that a conforming compiler must accept. You must be aware of these when porting applications between compilers. Table F-1 gives a summary of these limits.

In this table, a limit of `memory` indicates that the ARM compiler imposes no limit, other than that imposed by the available memory.

Table F-1 Implementation limits

Description	Recommended	ARM
Nesting levels of compound statements, iteration control structures, and selection control structures.	256	memory
Nesting levels of conditional inclusion.	256	memory
Pointer, array, and function declarators (in any combination) modifying an arithmetic, structure, union, or incomplete type in a declaration.	256	memory
Nesting levels of parenthesized expressions within a full expression.	256	memory
Number of initial characters in an internal identifier or macro name.	1 024	memory
Number of initial characters in an external identifier.	1 024	memory
External identifiers in one translation unit.	65 536	memory
Identifiers with block scope declared in one block.	1 024	memory
Macro identifiers simultaneously defined in one translation unit.	65 536	memory
Parameters in one function declaration.	256	memory
Arguments in one function call.	256	memory
Parameters in one macro definition.	256	memory
Arguments in one macro invocation.	256	memory
Characters in one logical source line.	65 536	memory
Characters in a character string literal or wide string literal after concatenation.	65 536	memory
Size of a C or C++ object (including arrays).	262 144	4 294 967 296
Nesting levels of #include file.	256	memory
Case labels for a **`switch`** statement, excluding those for any nested **`switch`** statements.	16 384	memory
Data members in a single class, structure, or union.	16 384	memory
Enumeration constants in a single enumeration.	4 096	memory
Levels of nested class, structure, or union definitions in a single **`struct`** declaration-list.	256	memory
Functions registered by `atexit()`.	32	33
Direct and indirect base classes.	16 384	memory
Direct base classes for a single class.	1 024	memory

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table F-1 Implementation limits (continued)

Description	Recommended	ARM
Members declared in a single class.	4096	memory
Final overriding virtual functions in a class, accessible or not.	16384	memory
Direct and indirect virtual bases of a class.	1024	memory
Static members of a class.	1024	memory
Friend declarations in a class.	4096	memory
Access control declarations in a class.	4096	memory
Member initializers in a constructor definition.	6144	memory
Scope qualifications of one identifier.	256	memory
Nested external specifications.	1024	memory
Template arguments in a template declaration.	1024	memory
Recursively nested template instantiations.	17	memory
Handlers per try block.	256	memory
Throw specifications on a single function declaration.	256	memory

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

F.2 Limits for integral numbers

Table F-2 gives the ranges for integral numbers in ARM C and C++. The `Endpoint` column of the table gives the numerical value of the range endpoint. The `Hex value` column gives the bit pattern (in hexadecimal) that is interpreted as this value by the ARM compiler. These constants are defined in the `limits.h` include file.

When entering a constant, choose the size and sign with care. Constants are interpreted differently in decimal and hexadecimal/octal. See the appropriate C or C++ standard, or any of the recommended C and C++ textbooks for more information, as described in *Further reading on page 2-26* of *Migration and Compatibility*.

Table F-2 Integer ranges

Constant	Meaning	Value	Hex value
`CHAR_MAX`	Maximum value of **`char`**	255	`0xFF`
`CHAR_MIN`	Minimum value of **`char`**	0	`0x00`
`SCHAR_MAX`	Maximum value of **`signed char`**	127	`0x7F`
`SCHAR_MIN`	Minimum value of **`signed char`**	–128	`0x80`
`UCHAR_MAX`	Maximum value of **`unsigned char`**	255	`0xFF`
`SHRT_MAX`	Maximum value of **`short`**	32767	`0x7FFF`
`SHRT_MIN`	Minimum value of **`short`**	–32768	`0x8000`
`USHRT_MAX`	Maximum value of **`unsigned short`**	65535	`0xFFFF`
`INT_MAX`	Maximum value of **`int`**	2147483647	`0x7FFFFFFF`
`INT_MIN`	Minimum value of **`int`**	–2147483648	`0x80000000`
`LONG_MAX`	Maximum value of **`long`**	2147483647	`0x7FFFFFFF`
`LONG_MIN`	Minimum value of **`long`**	–2147483648	`0x80000000`
`ULONG_MAX`	Maximum value of **`unsigned long`**	4294967295	`0xFFFFFFFF`
`LLONG_MAX`	Maximum value of **`long long`**	9.2E+18	`0x7FFFFFFF FFFFFFFF`
`LLONG_MIN`	Minimum value of **`long long`**	–9.2E+18	`0x80000000 00000000`
`ULLONG_MAX`	Maximum value of **`unsigned long long`**	1.8E+19	`0xFFFFFFFF FFFFFFFF`

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

F.3 Limits for floating-point numbers

This section describes the characteristics of floating-point numbers.

Table F-3 gives the characteristics, ranges, and limits for floating-point numbers. These constants are defined in the `float.h` include file.

Table F-3 Floating-point limits

Constant	Meaning	Value
`FLT_MAX`	Maximum value of **`float`**	3.40282347e+38F
`FLT_MIN`	Minimum normalized positive floating-point number value of **`float`**	1.175494351e–38F
`DBL_MAX`	Maximum value of **`double`**	1.79769313486231571e+308
`DBL_MIN`	Minimum normalized positive floating-point number value of **`double`**	2.22507385850720138e–308
`LDBL_MAX`	Maximum value of **`long double`**	1.79769313486231571e+308
`LDBL_MIN`	Minimum normalized positive floating-point number value of **`long double`**	2.22507385850720138e–308
`FLT_MAX_EXP`	Maximum value of base 2 exponent for type **`float`**	128
`FLT_MIN_EXP`	Minimum value of base 2 exponent for type **`float`**	–125
`DBL_MAX_EXP`	Maximum value of base 2 exponent for type **`double`**	1024
`DBL_MIN_EXP`	Minimum value of base 2 exponent for type **`double`**	–1021
`LDBL_MAX_EXP`	Maximum value of base 2 exponent for type **`long double`**	1024
`LDBL_MIN_EXP`	Minimum value of base 2 exponent for type **`long double`**	–1021
`FLT_MAX_10_EXP`	Maximum value of base 10 exponent for type **`float`**	38
`FLT_MIN_10_EXP`	Minimum value of base 10 exponent for type **`float`**	–37
`DBL_MAX_10_EXP`	Maximum value of base 10 exponent for type **`double`**	308
`DBL_MIN_10_EXP`	Minimum value of base 10 exponent for type **`double`**	–307
`LDBL_MAX_10_EXP`	Maximum value of base 10 exponent for type **`long double`**	308
`LDBL_MIN_10_EXP`	Minimum value of base 10 exponent for type **`long double`**	–307

Table F-4 describes other characteristics of floating-point numbers. These constants are also defined in the `float.h` include file.

Table F-4 Other floating-point characteristics

Constant	Meaning	Value
`FLT_RADIX`	Base (radix) of the ARM floating-point number representation	2
`FLT_ROUNDS`	Rounding mode for floating-point numbers	(nearest) 1
`FLT_DIG`	Decimal digits of precision for **`float`**	6
`DBL_DIG`	Decimal digits of precision for **`double`**	15

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential

Table F-4 Other floating-point characteristics (continued)

Constant	Meaning	Value
LDBL_DIG	Decimal digits of precision for **long double**	15
FLT_MANT_DIG	Binary digits of precision for type **float**	24
DBL_MANT_DIG	Binary digits of precision for type **double**	53
LDBL_MANT_DIG	Binary digits of precision for type **long double**	53
FLT_EPSILON	Smallest positive value of x that 1.0 + x != 1.0 for type **float**	1.19209290e–7F
DBL_EPSILON	Smallest positive value of x that 1.0 + x != 1.0 for type **double**	2.2204460492503131e–16
LDBL_EPSILON	Smallest positive value of x that 1.0 + x != 1.0 for type **long double**	2.2204460492503131e–16L

Note

- When a floating-point number is converted to a shorter floating-point number, it is rounded to the nearest representable number.
- Floating-point arithmetic conforms to IEEE 754.

Copyright © 2007-2008, 2011 ARM. All rights reserved.
Non-Confidential